

# Classical command of quantum systems via rigidity of CHSH games

Ben W. Reichardt

University of Southern California

Falk Unger

Knight Capital Group

Umesh Vazirani

UC Berkeley

Can a classical system command a general adversarial quantum system to realize arbitrary quantum dynamics? If so, then we could realize the dream of device-independent quantum cryptography: using untrusted quantum devices to establish a shared random key, with security based on the correctness of quantum mechanics. It would also allow for testing whether a claimed quantum computer is truly quantum. Here we report a technique by which a classical system can certify the joint, entangled state of a bipartite quantum system, as well as command the application of specific operators on each subsystem. This is accomplished by showing a strong converse to Tsirelson's optimality result for the Clauser-Horne-Shimony-Holt (CHSH) game: the only way to win many games is if the bipartite state is close to the tensor product of EPR states, and the measurements are the optimal CHSH measurements on successive qubits. This leads directly to a scheme for device-independent quantum key distribution. Control over the state and operators can also be leveraged to create more elaborate protocols for reliably realizing general quantum circuits.

Do the laws of quantum mechanics place any limits on how well a classical experimentalist can characterize the state and dynamics of a large quantum system? As a thought experiment, consider that we are presented with a quantum system, together with instructions on how to control its evolution from a claimed initial state. Can we, as classical beings, possibly convince ourselves that the quantum system was indeed initialized as claimed, and that its state evolves as we instruct?


More formally, model the quantum system as contained in a black box, and model our classical interactions with it as questions and answers across a digital interface, perhaps of buttons and light bulbs (Figure 1). Using this limited interface, we wish to characterize the initial state of the system. We also wish to certify that on command—by pressing a suitable sequence of buttons—the system applies a chosen local Hamiltonian, or equivalently a sequence of one- and two-qubit quantum gates, and outputs desired measurement results.

Although a philosophical question, a positive resolution would have important consequences. It is particularly relevant in quantum cryptography, where it is natural to

model the quantum system as adversarial since the goal is to protect honest users from malicious adversaries. The *raison d'être* of quantum cryptography is to create a cryptographic system with security premised solely on basic laws of physics, and with quantum key distribution (QKD) and its security proofs<sup>1–3</sup> it appeared to have achieved exactly this. However, attackers have repeatedly breached the security of QKD experiments, by exploiting imperfect implementations of the quantum devices.<sup>4–6</sup> Rather than relying on ad hoc countermeasures, Mayers and Yao's 1998 vision of device-independent (DI) QKD,<sup>7</sup> hinted at earlier by Ekert,<sup>8</sup> relaxes all modeling assumptions on the devices, and even allows for them to have been constructed by an adversary. It instead imagines giving the devices tests that cannot be passed unless they carry out the QKD protocol securely. The challenge at the heart of this vision is for a classical experimentalist to force untrusted quantum devices to act according to certain specifications. DIQKD has not been known to be possible; security proofs to date require the unrealistic assumption that the devices have no memory between trials, or that each party has many, strictly isolated devices.<sup>9–18</sup> A scheme for characterizing and commanding a black-box quantum device would provide a novel approach to achieving DIQKD.

Further, as the power of quantum mechanics is harnessed at larger scales, for example with the advent of quantum computers, it will be useful to evaluate whether a quantum device in fact carries out the claimed dynamics.<sup>19,20</sup> Finally, we might wish to test the applicability of quantum mechanics for large systems, a situation in which Nature itself plays the role of the adversary.<sup>21</sup>

The existence of a general scheme for commanding an adversarial quantum device appears singularly implausible. For example, in an adversarial setting, experiments cannot be repeated exactly to gather statistics, since a system with memory could deliberately deceive the experimentalist.



**Figure 1: Classical interaction with a quantum system.** A general system can be abstracted as a black box, with two buttons for accepting binary input and two light bulbs for output. Using this interface, we wish to control fully the system's quantum dynamics.




Figure 2: **Test for quantumness.** In a CHSH experiment, or “game,” the experimentalist Eve sends independent, uniformly random bits  $A$  and  $B$  to the devices Alice and Bob, respectively, who respond with bits  $X$  and  $Y$ . The devices “win” the game if  $AB = X \oplus Y$ . By a Bell inequality, classical devices can win with probability at most  $3/4$ . Quantum devices can win with probability  $\omega^* = \cos^2(\pi/8) \approx 85.4\%$ , if they follow an ideal CHSH strategy: on a shared Einstein-Podolsky-Rosen (EPR) state  $|\varphi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ , Bob measures the Pauli operator  $\sigma_z$  if  $B = 0$  or  $\sigma_x$  if  $B = 1$ , and Alice measures  $\frac{1}{\sqrt{2}}(\sigma_z + (-1)^A \sigma_x)$ . Tsirelson showed that this strategy is optimal.<sup>22</sup>

More fundamentally, as macroscopic, classical entities, our access to a quantum system is extremely limited and indirect, and the measurements we apply collapse the quantum state. Furthermore, whereas the dimension of the underlying Hilbert space scales exponentially in the number of particles or can be infinite, the information accessible via measurement only grows linearly. Indeed, as formulated it is impossible to command a single black-box system. Quite simply, one cannot distinguish between a quantum system that evolves as desired and a device that merely simulates the desired evolution using a classical computer.

In this paper, we consider a closely related scenario. Suppose we are instead given two devices, each modeled as a black box as above, and prevented from communicating with each other. In this setting, with no further assumptions, we show how to classically command the devices. That is, there is a strategy for pushing the buttons such that the answering light bulb flashes will satisfy a prescribed test only if the two devices started in a particular initial quantum state, to which they applied a desired sequence of quantum gates. Moreover, though impractical, the scheme is theoretically efficient—in the sense that the total effort, measured by the number of button pushes, scales as a polynomial function of the size of the desired quantum circuit. A DIQKD scheme follows, among other consequences.

The starting point for our protocol is the famous Bell experiment,<sup>23</sup> and its subsequent distillation by Clauser, Horne, Shimony and Holt (CHSH).<sup>24</sup> Conceptually modeled as a game (Figure 2), it provides a “test for quantumness,” a way for an experimentalist, whom we shall call Eve, to demonstrate the entanglement of two space-like separated devices, Alice and Bob.

Consider a protocol in which Eve plays a long sequence of CHSH games with Alice and Bob, and tests that they win close to the optimal fraction  $\omega^*$  of the games. This paper’s

main technical result establishes that if the devices pass Eve’s test with high probability, then at the beginning of a randomly chosen long subsequence of games, Alice and Bob must share many EPR states in tensor product, that they measure one at a time using the single-game ideal CHSH operators of Figure 2. This is a step towards the general vision outlined above because it characterizes the initial state of many qubits, and allows Eve to command the devices to perform certain single-qubit operations. Of course, we cannot hope to characterize the devices’ strategies exactly, but only for a suitable notion of approximation.

In order to make a more precise statement, first consider a single CHSH game. We show that if the devices win with probability  $\omega^* - \epsilon$ , then they must share a state that is  $O(\sqrt{\epsilon})$ -close to an EPR state, possibly in tensor product with an additional state. Moreover their joint measurement strategy is necessarily  $O(\sqrt{\epsilon})$ -close to the ideal strategy from Figure 2. (That is, applying Alice’s actual measurement operator to the shared state gets within distance  $O(\sqrt{\epsilon})$  of the result from applying her ideal measurement operator to the EPR state tensored with the ancilla; and similarly for Bob.) Since each device can locate its qubit share of the EPR state arbitrarily within its Hilbert space, these statements hold only up to local isometries. This may be seen as a robust converse to Tsirelson’s inequality, and as a rigidity property of the CHSH game: a nearly maximal Bell inequality violation rigidly locks into place the devices’ shared state and measurement directions.

A converse to Tsirelson’s inequality for the CHSH game has been shown previously in the exact case.<sup>25,26</sup> Robustness is important for applications, however, because the success probability of a system can never be known exactly. Robust,  $\epsilon > 0$ , converse statements have been shown based on a conjecture<sup>27</sup> or under restrictive symmetry assumptions.<sup>14,15</sup> Recently, robustness has independently been shown for the CHSH game.<sup>28,29</sup>

Scaling up to a sequence of  $n$  CHSH games, suppose Alice and Bob use a strategy such that they win at least  $(1 - \epsilon)\omega^*n$  of the games with high probability. By basic statistics, their strategy at the beginning of most games will win that single game with probability at least  $(1 - \epsilon^{\Omega(1)})\omega^*$ . Rigidity for the one-shot game therefore applies. However, their strategy for playing the  $j$ th game could depend on the previous games. The states close to EPR states used in different games could overlap significantly, and their locations could depend on the history. The multi-game rigidity theorem rules out such wayward behavior. It says that for most random blocks of  $m = n^{\Omega(1)}$  consecutive games, at the start of the block Alice and Bob must share a state that is close to a tensor product of  $m$  EPR states, tensored with an additional state, and must play each  $j$ th game by making measurements that are close to the ideal CHSH strategy on the  $j$ th EPR state—different games being entirely independent.

One way to view this theorem is that it scales up the

CHSH test for quantumness and allows for identifying many qubits' worth of entanglement. Much more than that, however, the multi-game rigidity theorem gives strong control over the devices' measurement operators for different games. As described below, combining the CHSH game protocol with protocols for state and process tomography, and for computation by teleportation, thereby gives a method for realizing arbitrary dynamics in quantum systems without making assumptions about the internal structure or operations. The dynamics are realized as the joint evolution of two isolated quantum systems, Alice and Bob, mediated by a classical experimentalist, Eve.

The problem of controlling computationally powerful but untrusted resources lies at the foundation of computer science. In the complexity class NP, for example, a polynomial-time routine—the “verifier”—is allowed one round of interaction with an arbitrarily powerful, but malicious, “prover.” We show that the same verifier can exploit the power of quantum-mechanical provers. In particular:

1. A classical verifier can efficiently simulate a quantum computer by interacting with two untrusted, polynomial-time quantum provers. This delegated computation scheme is also *blind*, meaning that each prover learns no more than the length of the computation.

2. The verifier in any quantum multi-prover interactive proof system can be assumed to be classical. Formally, the complexity classes QMIP and MIP\* are equal.

Previous work has considered a verifier who can store and control a constant number of qubits while interacting with a single prover.<sup>19,20,30,31</sup> Our work is also inspired by a proposal<sup>32</sup> that QMIP should equal MIP\*. Our protocol has a very different form, based on the multi-game rigidity theorem.

### Tensor-product structure for repeated games

A strategy  $\mathcal{S}$  for playing  $n$  sequential CHSH games specifies Alice and Bob's initial joint state as well as their measurement operators for every possible situation. That is, for  $X \in \{A, B\}$  and each  $j = 1, \dots, n$ ,  $\mathcal{S}$  specifies the measurement operators used by device  $X$  in game  $(j, h_{j-1}^D)$ , where  $h_{j-1}^D$  is any transcript of the device's input and output bits for the first  $j - 1$  games. For two strategies to be “close” means that the distributions of game transcripts they induce should be close in total variation distance; and that for most transcripts (drawn from either distribution), the resulting quantum states should be close in a suitable norm. We combine these conditions into one by defining for any strategy a block-diagonal density matrix that stores both the classical transcript and the resulting quantum state:

$$\rho_j = \bigoplus_{h_{j-1}} \Pr[h_{j-1}] \rho_j(h_{j-1}) . \quad (1)$$

Here  $h_{j-1} = (h_{j-1}^A, h_{j-1}^B)$  is the full transcript for the first  $j - 1$  games and  $\rho_j(h_{j-1})$  is the state at the beginning of

game  $j$  conditioned on  $h_{j-1}$ . Two strategies  $\mathcal{S}$  and  $\tilde{\mathcal{S}}$  are close if the associated  $\rho_j$  and  $\tilde{\rho}_j$  are close in trace distance, for every  $j$ .

Assume that for every  $j$  and most  $h_{j-1}$ , the devices' conditional joint strategy at the beginning of game  $j$  is “ $\epsilon$ -structured,” meaning that it wins with probability at least  $\omega^* - \epsilon$ . Our key theorem establishes that up to local basis changes, the devices' initial state must be close to  $n$  EPR states, possibly in tensor product with an irrelevant extra state, and that their total strategy  $\mathcal{S}$  must be close to an ideal strategy  $\hat{\mathcal{S}}$  that plays game  $j$  using the  $j$ th EPR state. Since the structure assumption can be established by martingale arguments on  $\text{poly}(n)$  sequential CHSH games, this implies the multi-game rigidity theorem.

The main challenge is to “locate” the ideal strategy  $\hat{\mathcal{S}}$  within Alice and Bob's Hilbert space, i.e., to find an isometry on each of their spaces under which their states and measurement operators are close to ideal. However, a priori, we do not know whether  $\mathcal{S}$  calls for the devices to measure actual qubits in each step, or even if so whether the qubits form EPR states, qubits for different games overlap each other, or the locations of the qubits depend on the outcomes of previous games.

The construction to locate the qubits in the single-game rigidity theorem is a good place to start. Consider an  $\epsilon$ -structured strategy, consisting of some shared mixed state in  $\mathcal{H}_A \otimes \mathcal{H}_B$ , and two-outcome projective measurements for each of Eve's possible questions. Truncate the devices' Hilbert spaces to finitely many dimensions, then decompose each space by Jordan's lemma<sup>33</sup> into the direct sum of two-dimensional spaces invariant under the projections. Within each such two-dimensional subspace adjust the angles between the projections to match the ideal strategy. The resulting operators define underlying qubits.<sup>1</sup>

For multiple CHSH games, the given strategy  $\mathcal{S}$  can be transformed into a nearby ideal strategy  $\hat{\mathcal{S}}$  by a three-step sequence:

1. First, replace each device's measurement operators by the ideal operators promised by the single-game rigidity theorem. In the resulting strategy  $\tilde{\mathcal{S}}$ , each device  $X$  plays every game  $(j, h_{j-1}^D)$  using the ideal CHSH game operators on some qubit, up to a local change in basis. However, the basis change can depend arbitrarily on  $h_{j-1}^D$ , and the qubits for different  $j$  need not be in tensor product.

2. In a “multi-qubit ideal strategy”  $\hat{\mathcal{S}}$ , the qubits used in each game can still depend on the local transcripts but must at least lie in tensor product with the qubits from previous games. This imposes a tensor-product subsystem structure that previous DIQKD proofs have assumed. The tensor-product structure is constructed beginning with a trivial transformation on  $\tilde{\mathcal{S}}$ : to each device, add  $n$  ancilla qubits each in state  $|0\rangle$ . Next, after a qubit has been measured, say as  $|\alpha_j\rangle$  in game  $j$ , swap it with the  $j$ th ancilla qubit, then rotate this fresh qubit from  $|0\rangle$  to  $|\alpha_j\rangle$

<sup>1</sup>A technical account with full proofs will appear elsewhere.<sup>34</sup>

and continue playing games  $j + 1, \dots, n$ . This defines a unitary change of basis that places the outcomes for games 1 to  $j$  in the first  $j$  ancilla qubits, and leaves the state in the original Hilbert space unchanged. Since qubits are set aside after being measured, the qubits for later games are automatically in tensor product with those for earlier games; the resulting strategy  $\tilde{\mathcal{S}}$  is multi-qubit ideal. At the end of the  $n$  games, swap back the ancilla qubits and undo their rotations, using the transcript.

3. In the last step, we replace  $\tilde{\mathcal{S}}$  with an ideal strategy  $\hat{\mathcal{S}}$ , in which Alice and Bob each play using a fixed set of  $n$  qubits. Fix a transcript  $\hat{h}_n$ , chosen at random. For the first time, change the devices' initial state: replace  $\rho_1$  with  $\hat{\rho}_1$ , a state having  $n$  EPR states in the locations determined by  $\hat{h}_n$  in  $\tilde{\mathcal{S}}$ . In  $\hat{\mathcal{S}}$ , the devices play using these EPR states, regardless of the actual transcript. This  $\hat{\mathcal{S}}$  is the desired ideal strategy.

### Ideal strategy $\hat{\mathcal{S}}$ is close to $\mathcal{S}$

It remains to show that the transformation's three steps incur a small error:  $\hat{\mathcal{S}}$  is close to  $\mathcal{S}$ . A major theme in the analysis is to leverage the known tensor-product structure between  $\mathcal{H}_A$  and  $\mathcal{H}_B$  to extract a tensor-product structure within  $\mathcal{H}_A$  and  $\mathcal{H}_B$ .

1.  $\mathcal{S} \approx \tilde{\mathcal{S}}$ : Although elementary, explaining this step is useful for establishing some notation. Let  $\rho_1$  be the devices' initial shared state, possibly entangled with the environment. Let  $\mathcal{E}_j^A$  and  $\mathcal{E}_j^B$  be the super-operators that implement Alice and Bob's respective strategies for game  $j$ ,  $\mathcal{E}_j^{AB} = \mathcal{E}_j^A \otimes \mathcal{E}_j^B$  and  $\mathcal{E}_{j,k}^{AB} = \mathcal{E}_k^{AB} \dots \mathcal{E}_j^{AB}$  for  $j \leq k$ ; thus the state  $\rho_j$  of Eq. (1) equals  $\mathcal{E}_{1,j-1}^{AB}(\rho_1)$ . For  $D \in \{A, B\}$ , let  $\tilde{\mathcal{E}}_j^D$  be the super-operator that replaces the actual measurement operators with the ideal operators promised by the CHSH rigidity theorem.  $\tilde{\mathcal{S}}$  is given by  $\rho_1$ ,  $\{\tilde{\mathcal{E}}_j^A\}$  and  $\{\tilde{\mathcal{E}}_j^B\}$ . If  $\Pr[\text{game } j \text{ is } \epsilon\text{-structured}] \geq 1 - \delta$ , then  $\|\mathcal{E}_j^{AB}(\rho_j) - \tilde{\mathcal{E}}_j^{AB}(\rho_j)\|_{\text{tr}} \leq 2\delta + O(\sqrt{\epsilon})$ . (This expression uses Eq. (1) to combine bounds on the probability of the bad event and the  $O(\sqrt{\epsilon})$  error from the good event.) To show our goal, that  $\mathcal{E}_{1,n}^{AB}(\rho_1) \approx \tilde{\mathcal{E}}_{1,n}^{AB}(\rho_1)$  in trace distance, use a hybrid argument that works backwards from game  $n$  to game 1 fixing each game's measurement operators one at a time.

2.  $\tilde{\mathcal{S}} \approx \hat{\mathcal{S}}$ : The key to showing that  $\tilde{\mathcal{S}}$  is close to  $\hat{\mathcal{S}}$  is the fact that operations on one half of an EPR state can equivalently be performed on the other half, since for any  $2 \times 2$  matrix  $M$ ,  $(M \otimes I)(|00\rangle + |11\rangle) = (I \otimes M^T)(|00\rangle + |11\rangle)$ . This means that the outcome of an  $\epsilon$ -structured CHSH game would be nearly unchanged if Bob were hypothetically to perform Alice's measurement before his own. By moving Alice's measurement operators for games  $j + 1$  to  $n$  over to Bob's side, we see that they cannot significantly affect the qubit  $|\alpha_j\rangle$  from game  $j$  on her side. Therefore, undoing the original change of basis restores the ancilla qubits nearly to their initial state  $|0^n\rangle$ , and  $\tilde{\mathcal{S}} \approx \hat{\mathcal{S}}$ .

Formally, define a unitary super-operator  $\mathcal{V}_j$  that rotates the  $j$ th ancilla qubit to  $|\alpha_j\rangle$ , depending on Alice's transcript  $h_j^A$ . Define a unitary super-operator  $\mathcal{T}_j$  to apply  $\mathcal{V}_j$  and swap the  $j$ th ancilla qubit with the qubit Alice uses in game  $j$  (depending on  $h_{j-1}^A$ ). Alice's multi-qubit ideal strategy is given by

$$\tilde{\mathcal{E}}_j^A = \mathcal{T}_{1,j-1}^{-1}(\mathbf{1}_{\mathbf{C}^{2^n}} \otimes \tilde{\mathcal{E}}_j^A)\mathcal{T}_{1,j-1} \quad (2)$$

We aim to show that the strategy given by  $\rho_1$ ,  $\{\tilde{\mathcal{E}}_j^A\}$  and  $\{\tilde{\mathcal{E}}_j^B\}$  is close to  $\tilde{\mathcal{S}}$  up to the fixed isometry that prepends  $|0^n\rangle\langle 0^n|$  to the state, i.e., that  $|0^n\rangle\langle 0^n| \otimes \tilde{\mathcal{E}}_{1,n}^{AB}(\rho_1) \approx \tilde{\mathcal{E}}_{1,n}^A(|0^n\rangle\langle 0^n| \otimes \tilde{\mathcal{E}}_{1,n}^B(\rho_1))$ . Define a super-operator  $\tilde{\mathcal{F}}_j^{AB}$ , in which Alice's measurements are made on Bob's Hilbert space  $\mathcal{H}_B$ , on the qubit determined by Bob's transcript  $h_{j-1}^B$ . Since most games are  $\epsilon$ -structured, by the CHSH rigidity theorem,  $\tilde{\mathcal{F}}_{j+1,k}^{AB}(\tilde{\rho}_{j+1}) \approx \tilde{\mathcal{E}}_{j+1,k}^{AB}(\tilde{\rho}_{j+1}) = \tilde{\rho}_{k+1}$  for  $j \leq k$ . Since  $\tilde{\mathcal{F}}_{j+1,k}^{AB}$  acts on  $\mathcal{H}_B$ , it does not affect Alice's qubit  $|\alpha_j\rangle$  from game  $j$  at all, and so this qubit must stay near  $|\alpha_j\rangle$  in  $\tilde{\rho}_{k+1}$  as well, i.e., the trace of the reduced density matrix against the projection  $|\alpha_j\rangle\langle \alpha_j|$  stays close to one. As this holds for every  $j$ ,  $\mathcal{T}_{1,n}^{-1}$  indeed returns the ancillas almost to their initial state  $|0^n\rangle$ . The  $\{\tilde{\mathcal{E}}_j^B\}$  are symmetrically adjusted to  $\{\tilde{\mathcal{E}}_j^A\}$ .

3.  $\tilde{\mathcal{S}} \approx \hat{\mathcal{S}}$ : In  $\tilde{\mathcal{S}}$ , Alice and Bob play according to a strategy in which every game uses a qubit in tensor product with the previous games' qubits. However, the qubit's location can depend on previous games' outcomes. We wish to argue that Alice and Bob must play using a single set of  $n$  qubits, fixed in advance independent of the transcript.

Intuitively, if the location of Alice's  $j$ th qubit depended on  $h_{j-1}^A$ , then since the devices cannot communicate with each other, Bob could not know which of his qubits to measure. However, Alice and Bob's transcripts are significantly correlated, and we must show that they cannot use these correlations to coordinate dynamically the locations of their qubits.

For a toy example that illustrates the issue, consider two devices who play the first  $n - 1$  games honestly and which at the beginning of the last game share two EPR states,  $|\varphi\rangle^{\otimes 2}$ . Say that for certain functions  $f$  and  $g$ , Alice uses EPR state  $f(h_{n-1}^A) \in \{1, 2\}$  in game  $n$ , and Bob uses EPR state  $g(h_{n-1}^B) \in \{1, 2\}$ . For game  $n$  to be structured, they need  $f(h_{n-1}^A) = g(h_{n-1}^B)$  so that they measure the same EPR state. Now Alice and Bob's local transcripts are each uniformly random, separately, but they have a constant correlation in every game coordinate. It is straightforward to argue based on coordinate influence that if  $\Pr[f(H_{n-1}^A) \neq g(H_{n-1}^B)]$  is small, then  $f$  and  $g$  must both be nearly constant. Thus one of the two EPR states is used almost always.

This example gives an essentially classical cheating strategy. The actual devices may be significantly more sophisticated. In particular, small amounts of cheating in earlier



games might enable an avalanche of more and more blatant cheating in later games, drastically changing the underlying quantum state. If, for example, Alice knowingly manages to swap her halves of the two last EPR states along some transcripts  $h_{n-1}^A$ , then she can use completely different strategies for the last game without having to coordinate with Bob. We control such errors, like in the arguments sketched above, by approximating Alice's super-operator as acting on Bob's Hilbert space. More formal arguments are deferred to the supplementary material.

The conclusion that the devices' joint strategy is close to ideal,  $\mathcal{E}_{1,n}^{AB}(\rho_1) \approx \hat{\mathcal{E}}_{1,n}^{AB}(\hat{\rho}_1)$ , is not strong enough for our applications, in which sometimes Eve plays CHSH games with only one of the two devices. We need to show that the devices' strategies are *separately* close to ideal, i.e.,  $\mathcal{E}_{1,n}^A(\rho_1) \approx \hat{\mathcal{E}}_{1,n}^A(\hat{\rho}_1)$  and  $\mathcal{E}_{1,n}^B(\rho_1) \approx \hat{\mathcal{E}}_{1,n}^B(\hat{\rho}_1)$ . These estimates cannot be obtained directly because our main assumption, that every game  $j$  is usually  $\epsilon$ -structured, is only of use if both devices have played games 1 through  $j-1$ —it gives information about  $\mathcal{E}_j^D$  applied to  $\mathcal{E}_{1,j-1}^{AB}(\rho_1)$ , not about  $\mathcal{E}_j^D$  applied to  $\mathcal{E}_{1,j-1}^D(\rho_1)$ . The key idea to obtain separate estimates is that applying both devices' super-operators is almost equivalent to applying Alice's super-operator, *guessing* Bob's measurement outcome from the ideal conditional distribution, and based on the guess applying a controlled unitary correction to his qubit. Since Alice's super-operator collapses both qubits of the EPR state, it is not actually necessary to measure Bob's qubit. Defining  $\mathcal{G}_j^B$  to be this guess-and-correct super-operator, two hybrid arguments give  $\mathcal{E}_{1,n}^{AB}(\rho_1) \approx \mathcal{G}_{1,n}^B \mathcal{E}_{1,n}^A(\rho_1)$  and  $\tilde{\mathcal{E}}_{1,n}^A \mathcal{E}_{1,n}^B(\rho_1) \approx \mathcal{G}_{1,n}^B \tilde{\mathcal{E}}_{1,n}^A(\rho_1)$ . Thus,

$$\mathcal{G}_{1,n}^B \mathcal{E}_{1,n}^A(\rho_1) \approx \mathcal{G}_{1,n}^B \tilde{\mathcal{E}}_{1,n}^A(\rho_1) .$$

The same super-operator  $\mathcal{G}_{1,n}^B$  appears on both the left- and right-hand sides above. In general, applying a super-operator can reduce the trace distance. In this case, however, it does not; the correction part of  $\mathcal{G}_{1,n}^B$  is unitary, and the guessing part is a stochastic map acting on a copy of Alice's classical transcript register. Therefore, indeed  $\mathcal{E}_{1,n}^A(\rho_1) \approx \tilde{\mathcal{E}}_{1,n}^A(\rho_1)$ . The third step of the proof uses a similar, but more involved, argument.

## Verified quantum dynamics

Our scheme for verified quantum dynamics is based on the idea of computation by teleportation, which reduces computation to preparing certain resource states and applying Bell measurements (Figure 3(f)).<sup>35</sup> Say that Eve wants to simulate a quantum circuit  $\mathcal{C}$ , over the gate set  $\{H, G, \text{CNOT}\}$ , where  $H$  is the Hadamard gate and  $G = \exp(-i\frac{\pi}{8}\sigma_y)$ . Eve asks Bob to prepare many copies of the state  $|0\rangle \otimes (I \otimes H)|\varphi\rangle \otimes (I \otimes G)|\varphi\rangle \otimes \text{CNOT}_{2,4}(|\varphi\rangle \otimes |\varphi\rangle)$ . He can do so by applying one-, two- and four-qubit measurements to his halves of the shared EPR states and reporting

the results to Eve. If he plays honestly, Alice's shares of the EPR states collapse into the desired resource states, up to simple corrections. Each resource state corresponds to a basic operation in  $\mathcal{C}$ . Eve wires these up by repeatedly directing Alice to make a Bell measurement connecting the output of one operation to the input of the next operation in  $\mathcal{C}$ . After each  $G$  gate, an  $H$  correction might be required.

Of course, Alice and Bob might not follow directions. To enforce honest play, Eve runs this protocol only a small fraction of the time, and otherwise chooses uniformly between three alternative protocols sketched in Figure 3. Let  $m = |\mathcal{C}|^{O(1)}$  and  $n = m^{O(1)}$ .

1. In the "state tomography" protocol, Eve chooses  $K$  uniformly from  $\{1, \dots, n/m\}$ . She referees  $(K-1)m$  CHSH games. Then in the  $K$ th block of  $m$ , Eve asks Bob to prepare the resource states, in a random order, while continuing to play CHSH games with Alice. Eve rejects if the tomography statistics are inconsistent. We prove that if Alice plays honestly and Eve accepts with high probability, then on most randomly chosen small subsets of the resource state positions, Alice's reduced state is close to the correct tensor product of resource states.

2. In the "process tomography" protocol, Eve again chooses  $K$  uniformly from  $\{1, \dots, n/m\}$  and referees  $(K-1)m$  CHSH games. In the  $K$ th block of  $m$ , Eve asks Alice to make Bell measurements on random pairs of qubits, while continuing to play CHSH games with Bob. If Alice's reported result for *any* pair of qubits is inconsistent with Bob's outcomes, Eve rejects. Then if Bob plays honestly and Eve accepts with high probability, Alice must also have applied the Bell measurements honestly.

3. In the third protocol, Eve simply referees  $n$  sequential CHSH games with both devices and rejects if they do not win at least  $(1-\epsilon)\omega^*n$  games.

From Bob's perspective the process tomography and computation protocols are indistinguishable, as are the state tomography and CHSH game protocols. From Alice's perspective, the state tomography and computation protocols are indistinguishable, as are the process tomography and CHSH game protocols. The devices must behave identically in indistinguishable protocols. The multi-game rigidity theorem therefore provides the base for a chain of implications that implies that if Eve accepts with high probability, then the devices must implement  $\mathcal{C}$  honestly.

Four main technical problems obstruct these claims.

First, in the state tomography protocol, if Bob is dishonest, then Alice gets an arbitrary  $m$ -qubit state, and there is no reason why it should split into a tensor product of repeated, constant-qubit states. Nonetheless, we argue using martingales that if the counts of Alice's different measurement outcomes roughly match their expectations with high probability, then for most reported measurement outcomes from Bob and for most subsystems  $j$ , Alice's conditional state reduced to her  $j$ th subsystem is close to what it should be.

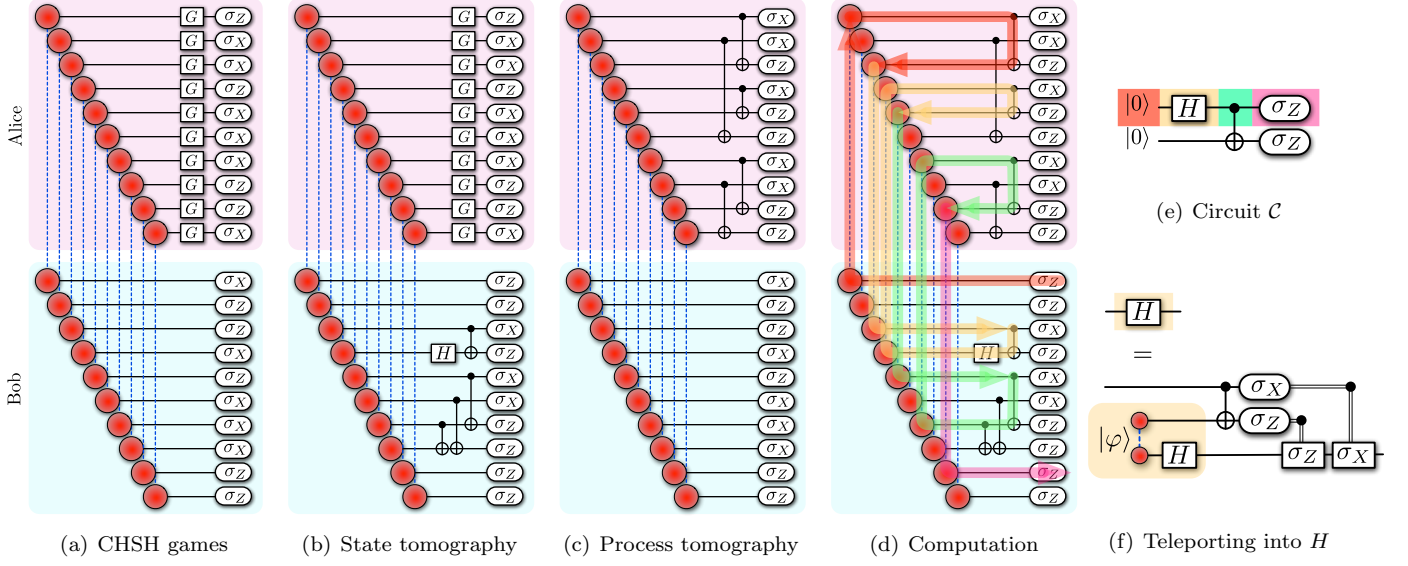


Figure 3: **Sub-protocols for verified quantum dynamics.** To delegate a quantum computation, Eve runs a random one of four sub-protocols with Alice and Bob. **a**, Playing many CHSH games ensures that the devices play honestly using shared EPR states. **b-c**, This lets Eve apply state or process tomography to characterize more complicated multi-qubit operations. **d-e**, By adaptively combining these operations, Eve directs a quantum circuit  $\mathcal{C}$ . The zig-zagging logical path of the first qubit of  $\mathcal{C}$  is highlighted. **f**, Each gate of  $\mathcal{C}$  is implemented through teleportation; in this simpler example,  $H$  is applied by a Bell measurement on half of the resource state  $(I \otimes H)|\varphi\rangle$ .

Furthermore, saturating Tsirelson’s inequality for the CHSH game only implies that Alice is honestly making Pauli  $\sigma_x$  and  $\sigma_z$  measurements on her half of an EPR state. Tomography also requires  $\sigma_y$  measurements. To sidestep this issue, we generalize a theory introduced by McKague<sup>36</sup> and prove that there is a large class of states, including the necessary resource states, that are all robustly determined by only  $\sigma_x$  and  $\sigma_z$  measurements.

A bigger problem, though, is that we want to characterize the operations that the devices apply to their shared EPR states, and not just the states that these operations create on the other side. The distinction is the same as that between process and state tomography. Essentially, the problem is that the correct states could be generated by incorrect processes. Moreover, as for sequential CHSH games, Bob’s strategy in early tomography rounds might be sufficiently dishonest as to allow him in later rounds to apply completely dishonest operators. A key observation to avoid this problem is that it is enough to certify the states prepared by one device and the processes applied by the other. Then since a broad class of states can be certified, for applications it suffices to certify a much smaller set of operations. We restrict consideration to Pauli stabilizer measurements.<sup>37</sup> For Pauli operators in the stabilizer of a state, the measurement outcome is deterministic. Therefore if Alice reports the wrong stabilizer syndrome in even a single round, Eve can reject. Our process certification analysis is similar to some of the arguments used above. We argue that Alice’s earlier measurements cannot usually

overly disturb the qubits intended for use in later measurements, by pulling Alice’s measurement super-operators over onto Bob’s halves of the EPR states.

Finally, the verifier’s questions in the state and process tomography protocols are non-adaptive, whereas in computation by teleportation the questions must be chosen adaptively based on previous responses. This is an attack vector in some related protocols. However, we argue that the devices can learn nothing from the adaptive questions. This follows because computation by teleportation can be implemented exactly equivalently either by choosing Bob’s state preparation questions non-adaptively and Alice’s process questions adaptively, or vice versa.

The proof that  $\text{QMIP} = \text{MIP}^*$  follows along similar lines. Begin with a  $k$ -prover protocol. We may assume that it has two rounds of quantum messages from the provers, before and after the verifier broadcasts a random bit.<sup>38</sup> To convert to an  $\text{MIP}^*$  protocol, with classical messages, add two additional provers, Alice and Bob. Eve teleports the original  $k$  provers’ messages to Alice, and directs Alice and Bob together to apply the quantum verifier’s acceptance predicate.

## Discussion

By characterizing the device strategies that can win many successive CHSH games, we have shown how a fully classical party can direct the actions of two untrusted quantum devices. The simplest case is device-independent quan-

tum key distribution, free of the independence assumptions needed in previous analyses. A major open problem is to improve the efficiency of our schemes, so as to get a key rate that is a positive constant instead of inverse polynomially small. More generally, the CHSH game rigidity theorems have no classical analog; and further implications for cryptographic primitives and protocols with drastically reduced security assumptions remain to be explored.

## A Rigidity for the CHSH game

In this appendix, we sketch the proof of the single-game CHSH rigidity theorem for the case that the devices' Hilbert spaces  $\mathcal{H}_A$  and  $\mathcal{H}_B$  are finite dimensional, and  $\epsilon = 0$ , i.e., the devices achieve the maximum correlation allowed by Tsirelson's inequality. Hilbert space completeness allows for truncating an infinite-dimensional space to finitely many dimensions, at an arbitrarily small cost.

A general strategy for the devices consists of measuring some shared state in  $\mathcal{H}_A \otimes \mathcal{H}_B$ . Since the success probability is extremal, i.e.,  $\epsilon = 0$ , we may assume that the state is extremal, i.e., is pure. Each device measures its system using a reflection that depends on Eve's question, and returns the sign of the observed eigenvalue  $\pm 1$ . The shared state  $|\psi\rangle$  and Alice and Bob's four reflections  $R_0^A, R_1^A, R_0^B$  and  $R_1^B$  determine the strategy.

Jordan's Lemma<sup>33</sup> states that any two reflections acting on a finite-dimensional space can be simultaneously block-diagonalized into  $1 \times 1$  and  $2 \times 2$  blocks. (The same statement is false for infinite-dimensional spaces.) Apply the lemma to  $R_0^A$  and  $R_1^A$ , to obtain

$$R_a^A = \bigoplus_i R_a^A(i), \quad (3)$$

where  $i$  labels the block index, and each  $R_a^A(i)$  is a  $1 \times 1$  or  $2 \times 2$  reflection. By adding placeholder dimensions, we may assume without loss of generality that each block is  $2 \times 2$ . Eq. (3), which can equivalently be rewritten  $R_a^A = \sum_i R_a^A(i) \otimes |i\rangle\langle i|$ , gives a basis in which  $\mathcal{H}_A = \mathbf{C}^2 \otimes \mathcal{H}'_A$ , where  $\mathcal{H}'_A$  is the Hilbert space with orthonormal basis  $\{|i\rangle\}$ . Thus Jordan's Lemma gives (an extension of) the a priori formless space  $\mathcal{H}_A$  a tensor-product structure. It locates within  $\mathcal{H}_A$  a qubit  $\mathbf{C}^2$ . However, Alice's operators need not act locally on this qubit; a controlled rotation is still needed to align her operators. After this rotation, we will show that Alice's strategy is close to the ideal CHSH game strategy that measures this qubit using the operators given in Figure 2.

Since the measurement of the block index commutes with both  $R_a^A$ , we may assume that Alice measures the block index first. Thus we reduce to the case that  $\mathcal{H}_A = \mathbf{C}^2$ , and, by a symmetrical argument, that  $\mathcal{H}_B = \mathbf{C}^2$ , still with  $\epsilon = 0$ .

If the  $R_a^D$  reflections act on  $\mathbf{C}^2$  and are not equal to  $\pm I$ , then we can choose a basis such that  $R_0^D = \sigma_z =$

$\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ ,  $R_1^A = \begin{pmatrix} \cos 2\theta & \sin 2\theta \\ \sin 2\theta & -\cos 2\theta \end{pmatrix}$  and  $R_1^B = \begin{pmatrix} \cos 2\theta' & \sin 2\theta' \\ \sin 2\theta' & -\cos 2\theta' \end{pmatrix}$  for certain angles  $\theta, \theta' \in [0, \frac{\pi}{2}]$ . (As this basis depends on the block index, changing basis amounts to a controlled qubit rotation.) Letting  $M_a = \frac{1}{2}(R_0^A + (-1)^a R_1^A) \otimes I - \frac{1}{\sqrt{2}}I \otimes R_a^B$  for  $a \in \{0, 1\}$ , the success probability satisfies

$$\begin{aligned} 2\sqrt{2} - 8\epsilon &\leq 8\Pr[AB = X \oplus Y] - 4 \\ &= \langle \psi | \left( \sum_{a,b \in \{0,1\}} (-1)^{ab} R_a^A \otimes R_b^B \right) | \psi \rangle \\ &= 2\sqrt{2} - \sqrt{2} \langle \psi | (M_0^2 + M_1^2) | \psi \rangle. \end{aligned}$$

For  $\epsilon = 0$ , this means that  $|\psi\rangle$  must lie in the intersection of the kernels of  $M_0$  and  $M_1$ . The four eigenvalues of  $M_0$  are  $\pm \cos \theta \pm \frac{1}{\sqrt{2}}$ . For the kernel to be nonempty, it must be that  $\theta = \frac{\pi}{4}$ . A symmetrical argument implies that  $\theta' = \frac{\pi}{4}$ . For small  $\epsilon > 0$ ,  $|\psi\rangle$  must lie close to small-eigenvalue subspaces of both  $M_0$  and  $M_1$ , implying that  $\theta$  and  $\theta'$  are close to  $\frac{\pi}{4}$ . Thus the measurement operators are rigidly determined.

For  $\theta = \theta' = \frac{\pi}{4}$ , the kernel of  $\sqrt{2}((HG) \otimes I)M_0((G^\dagger H) \otimes I) = \sigma_z \otimes I - I \otimes \sigma_z$  is spanned by the vectors  $|00\rangle$  and  $|11\rangle$ . The kernel of  $\sqrt{2}((HG) \otimes I)M_1((G^\dagger H) \otimes I) = \sigma_x \otimes I - I \otimes \sigma_x$  is spanned by the vectors  $|+\rangle \otimes |+\rangle = \frac{1}{2}(|00\rangle + |01\rangle + |10\rangle + |11\rangle)$  and  $|-\rangle \otimes |-\rangle = \frac{1}{2}(|00\rangle - |01\rangle - |10\rangle + |11\rangle)$ . For the  $|01\rangle$  and  $|10\rangle$  terms to cancel out, a linear combination of these vectors must have equal coefficients. The intersection between the two kernels is therefore spanned by  $|00\rangle + |11\rangle$ . Thus the state  $|\psi\rangle$  is rigidly determined.

The above argument conveys much of the intuition for the CHSH rigidity theorem. The case  $\epsilon > 0$  can be handled by maintaining suitable approximations. However, we have not explained the derivation of the operators  $M_0$  and  $M_1$ , chosen to satisfy  $\sum_{a,b \in \{0,1\}} (-1)^{ab} R_a^A \otimes R_b^B = 2\sqrt{2}I \otimes I - \sqrt{2}(M_0^2 + M_1^2)$ . In general, for a game in which Eve draws her questions from the distribution  $p(a, b)$  and accepts if  $x \oplus y = V(a, b)$ , let  $\Theta = \sum_{a,b} p(a, b) (-1)^{V(a,b)} |a\rangle\langle b|$  and  $\hat{\Theta} = \begin{pmatrix} 0 & \Theta \\ \Theta^\dagger & 0 \end{pmatrix}$ . Let  $\omega^*$  be the optimal success probability. By the Tsirelson semi-definite program,<sup>39</sup> the optimal bias is  $2\omega^* - 1 = \frac{1}{2} \max_{\Gamma \succeq 0, \Gamma \circ I = I} \langle \hat{\Theta}, \Gamma \rangle = \frac{1}{2} \min_{\Delta = \Delta \circ I \succeq \hat{\Theta}} \text{Tr } \Delta$ .  $\Gamma$  is the Gram matrix of the vectors  $R_a^A|\psi\rangle$  and  $R_b^B|\psi\rangle$ . Letting  $\Delta^*$  achieve the second optimum, we have  $\frac{1}{2} \langle \hat{\Theta}, \Gamma \rangle = (2\omega^* - 1) - \frac{1}{2} \langle \Delta^* - \hat{\Theta}, \Gamma \rangle$ . For the CHSH game,  $\Delta^* = \frac{1}{2\sqrt{2}}\mathbf{1}$ , and the matrices  $M_0, M_1$  correspond to eigenvectors of  $\Delta^* - \hat{\Theta}$ .

## B Rigidity for sequential CHSH games

For sequentially repeated CHSH games, our goal is to locate in  $\mathcal{H}_A$  and  $\mathcal{H}_B$  not one but many qubits, in tensor product, such that the devices' actual strategy is close to the ideal strategy that measures these qubits one at a time in sequence. In this section, we will introduce the notation and put together the theorems needed to prove rigidity for sequential CHSH games.

## B.1 Notation

To make our claims precise, we begin with some notation for CHSH games played in sequence, one following the next, with no communication between games.

A strategy  $\mathcal{S}$  for Alice and Bob to play  $n$  sequential CHSH games consists of the devices' Hilbert spaces, initial shared state and the reflections they use to play each game.

**Initial state:** Let  $\mathcal{H}_A$  and  $\mathcal{H}_B$  be Alice and Bob's respective Hilbert spaces, and  $\mathcal{H}_C$  any external space. Let  $|\psi\rangle \in \mathcal{H}_A \otimes \mathcal{H}_B \otimes \mathcal{H}_C$  be the devices' initial shared state.

**Transcripts:** Denote questions asked to Alice by  $a_1, \dots, a_n$ , questions asked to Bob by  $b_1, \dots, b_n$ , and possible answers by  $x_1, \dots, x_n$  and  $y_1, \dots, y_n$ , respectively. Write  $h_j^A = (a_1, x_1, \dots, a_j, x_j)$ ,  $h_j^B = (b_1, y_1, \dots, b_j, y_j)$  and  $h_j = (h_j^A, h_j^B)$ , a full transcript for games 1 through  $j$ . Write  $h_{j,k}$  and  $h_{j,k}^D$  for the full or partial transcripts for games  $j$  through  $k$ , inclusive.

**Reflections:** In game  $j$ , for questions  $a_j$  and  $b_j$ , let  $R_{a_j}^A(h_{j-1}^A)$  and  $R_{b_j}^B(h_{j-1}^B)$  be the reflections specifying Alice and Bob's respective strategies for game  $j$ , depending on the previous games' transcript. Define projections  $P_j^A(h_j^A) = \frac{1}{2}(\mathbf{1} + (-1)^{x_j} R_{a_j}^A(h_{j-1}^A))$  and  $P_j^B(h_j^B) = \frac{1}{2}(\mathbf{1} + (-1)^{y_j} R_{b_j}^B(h_{j-1}^B))$ . For  $D \in \{A, B\}$  and  $j \leq k$ , let  $P_{j,k}^D(h_k^D) = P_k^D(h_k^D) \cdots P_{j+1}^D(h_{j+1}^D) P_j^D(h_j^D)$ . Let  $P_{j,k}^{AB}(h_k) = P_{j,k}^A(h_k^A) \otimes P_{j,k}^B(h_k^B)$ .

**Super-operators:** For  $j < k$  and partial transcript  $h_j$ , define super-operators  $\mathcal{E}_k^{A|h_j^A}$  and  $\mathcal{E}_k^{B|h_j^B}$  by

$$\begin{aligned} \mathcal{E}_k^{D|h_j^D}(|h_{j+1,k-1}^D\rangle\langle h_{j+1,k-1}^D| \otimes \rho) \\ = \frac{1}{2} \sum_{\alpha_k, \chi_k} |h_{j+1,k}^D\rangle\langle h_{j+1,k}^D| \otimes P_k^D(h_k^D) \rho P_k^D(h_k^D), \end{aligned} \quad (4)$$

where  $h_k^D = (h_{k-1}^D, \alpha_k, \chi_k)$ . These super-operators capture the effects of Alice and Bob playing game  $k$ , where games  $j+1$  to  $k-1$  of the transcript are stored in a separate register. For  $\ell \geq k$ , let  $\mathcal{E}_{k,\ell}^{D|h_j^D} = \mathcal{E}_\ell^{D|h_j^D} \cdots \mathcal{E}_{k+1}^{D|h_j^D} \mathcal{E}_k^{D|h_j^D}$  and  $\mathcal{E}_{k,\ell}^{AB|h_j} = \mathcal{E}_{k,\ell}^{A|h_j^A} \otimes \mathcal{E}_{k,\ell}^{B|h_j^B}$ .

Let  $\rho_1 = |\psi\rangle\langle\psi|$ ,  $\rho_j(h_{j-1}) = P_{1,j-1}^{AB}(h_{j-1}) \rho_1 P_{1,j-1}^{AB}(h_{j-1})^\dagger / \text{Tr}(P_{1,j-1}^{AB}(h_{j-1}) \rho_1)$  be the state at the beginning of game  $j$  conditioned on  $h_{j-1}$ , and  $\rho_j = \mathcal{E}_{1,j-1}^{AB}(\rho_1) = \frac{1}{4^{j-1}} \sum_{h_{j-1}} |h_{j-1}\rangle\langle h_{j-1}| \otimes P_{1,j-1}^{AB}(h_{j-1}) \rho_1 P_{1,j-1}^{AB}(h_{j-1})^\dagger$ . Compare to Eq. (1) in the main text.

For fixed Hilbert spaces, a strategy  $\mathcal{S}$  can be identified with the tuple  $(\rho_1, \{\mathcal{E}_j^A\}, \{\mathcal{E}_j^B\})$ . When considering multiple strategies, we will decorate this notation to indicate the corresponding strategy, e.g.,  $\tilde{\mathcal{S}} = (\tilde{\rho}_1, \{\tilde{\mathcal{E}}_j^A\}, \{\tilde{\mathcal{E}}_j^B\})$ .

Call a strategy for a single CHSH game  $\epsilon$ -structured if the probability of winning is at least  $\omega^* - \epsilon/8$ . In our theorem, we will assume that most games the devices play are  $\epsilon$ -structured, in the following sense:

**Definition B.1** (Structured strategy). *A sequential CHSH game strategy  $\mathcal{S}$  is  $(\delta, \epsilon)$ -structured if for every  $j$ , there is at least a  $1 - \delta$  probability over transcripts  $h_{j-1}$  that game  $(j, h_{j-1})$  is  $\epsilon$ -structured.  $\mathcal{S}$  is  $\epsilon$ -structured if it is  $(\epsilon, \epsilon)$ -structured.*

We aim to show that the devices play nearly ideally:

**Definition B.2** (Ideal strategy). *A strategy  $\mathcal{S}$  for  $n$  sequential CHSH games is an ideal strategy if there exist isometries  $\mathcal{I}^D : \mathcal{H}_D \hookrightarrow (\mathbf{C}^2)^{\otimes n} \otimes \mathcal{H}'_D$  and a state  $|\psi'\rangle \in \mathcal{H}'_A \otimes \mathcal{H}'_B \otimes \mathcal{H}_C$  such that for every  $j$  and  $h_{j-1}$ ,*

$$\begin{aligned} \mathcal{I}^A \otimes \mathcal{I}^B |\psi\rangle &= |\varphi\rangle^{\otimes n} \otimes |\psi'\rangle \\ R_\alpha^D(h_{j-1}^D) &= \mathcal{I}^{D\dagger} (R_\alpha^D)_j \mathcal{I}^D, \end{aligned} \quad (5)$$

where  $(R_\alpha^D)_j$  denotes the ideal reflection operator  $R_\alpha^D$  from Figure 2 in the main text acting on the  $j$ th qubit.

In order to compare strategies, define a notion of simulation:

**Definition B.3** (Strategy simulation). *Let  $\mathcal{S}$  and  $\tilde{\mathcal{S}}$  be two strategies for playing  $n$  sequential CHSH games. For  $\epsilon \geq 0$ , we say that strategy  $\tilde{\mathcal{S}}$   $\epsilon$ -simulates strategy  $\mathcal{S}$  if they both use the same Hilbert spaces and for all  $j$ ,*

$$\max_{D \in \{A, B\}} \|\mathcal{E}_{1,j}^D(\rho_1) - \tilde{\mathcal{E}}_{1,j}^D(\tilde{\rho}_1)\|_{\text{tr}} \leq \epsilon. \quad (6)$$

Say that  $\tilde{\mathcal{S}}$  weakly  $\epsilon$ -simulates  $\mathcal{S}$  if only the weaker inequality  $\|\mathcal{E}_{1,j}^{AB}(\rho_1) - \tilde{\mathcal{E}}_{1,j}^{AB}(\tilde{\rho}_1)\|_{\text{tr}} \leq 2\epsilon$  holds.

It is also convenient to allow a basis change by local unitaries or local isometries:

**Definition B.4.** *A strategy  $\tilde{\mathcal{S}}$  is an isometric extension of  $\mathcal{S}$  if there exist isometries  $\mathcal{X}^D : \mathcal{H}_D \hookrightarrow \tilde{\mathcal{H}}_D$  such that  $|\psi\rangle = \mathcal{X}^A \otimes \mathcal{X}^B |\psi\rangle$  and  $\mathcal{X}^D R_\alpha^D(h_{j-1}^D) = \tilde{R}_\alpha^D(h_{j-1}^D) \mathcal{X}^D$  always. (Thus  $\mathcal{X}^D \mathcal{E}_j^D = \tilde{\mathcal{E}}_j^D \mathcal{X}^D$ .)*

## B.2 Main rigidity theorem and proof outline

Our main theorem states that a structured strategy can be closely simulated by an ideal strategy:

**Theorem B.5** (Rigidity theorem for sequential CHSH games). *There exists a constant  $\kappa$  such that for any  $\epsilon$ -structured strategy  $\mathcal{S}$  for  $n$  sequential CHSH games, there exists an ideal strategy  $\hat{\mathcal{S}}$  that  $\kappa n^\kappa \epsilon^{1/\kappa}$ -simulates an isometric extension of  $\mathcal{S}$ .*

The first step of the proof of Theorem B.5 is to replace the structured strategy  $\mathcal{S}$  with one in which the devices play every game using the ideal CHSH game operators on some qubit, up to a local change in basis. See Figure 4.



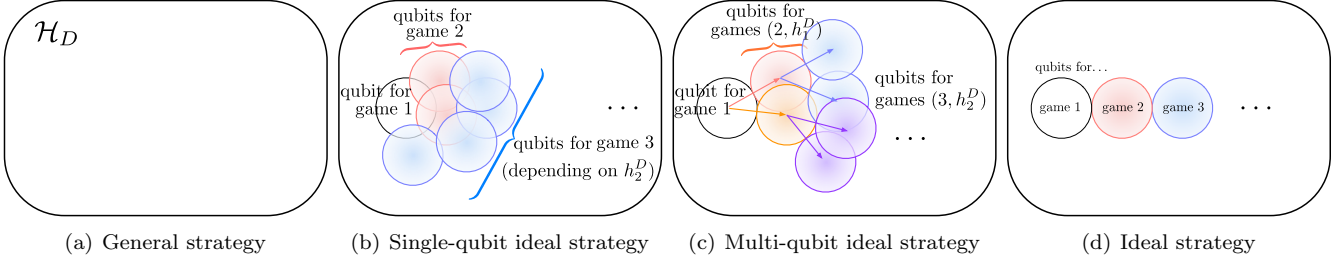


Figure 4: Proof outline for [Theorem B.5](#). **a**, Initially, each device  $D \in \{A, B\}$  can play arbitrarily, measuring in game  $j$  one of two reflections  $R_\alpha^D(h_{j-1}^D)$ ,  $\alpha \in \{0, 1\}$ , that depend on the local transcript  $h_{j-1}^D$  for the previous games. No structure is given for the Hilbert space  $\mathcal{H}_D$ . **b**, We first show that  $D$ 's strategy is close to a “single-qubit ideal strategy,” in which for every game it measures some qubit using the ideal CHSH game strategy, but the qubit locations can be arbitrary. Here, the qubits are illustrated schematically as balls, and the overlaps indicate that they need not be in tensor product. **c**, We then construct a nearby “multi-qubit ideal strategy,” in which the qubits used in each game must lie in tensor product with the qubits from previous games, but can overlap qubits used along other transcripts. **d**, Finally, we argue that the qubit locations cannot depend significantly on the transcript, and therefore that the original strategy is well-approximated by an ideal strategy that measures a fixed set of  $n$  qubits in sequence. (Note that these visualizations, representing qubits as balls, are inherently imprecise. A qubit's location in a Hilbert space is given not by a ball, but by the two anti-commuting reflection operators  $\sigma_x$  and  $\sigma_z$ .)

**Definition B.6** (Single-qubit ideal strategy). *A strategy  $\mathcal{S}$  is a single-qubit ideal strategy if there exist unitaries  $U_j^D(h_{j-1}^D) : \mathcal{H}_D \xrightarrow{\cong} \mathbf{C}^2 \otimes \mathcal{H}'_D$  such that always*

$$R_\alpha^D(h_{j-1}^D) = U_j^D(h_{j-1}^D)^\dagger (R_\alpha^D \otimes \mathbf{1}) U_j^D(h_{j-1}^D). \quad (7)$$

*That is, each device's reflections for game  $(j, h_{j-1}^D)$  are equivalent up to local unitaries to the ideal CHSH game reflections, although the qubits used need not be in tensor product.*

**Theorem B.7.** *There exists a constant  $\kappa$  such that if  $\mathcal{S}$  is an  $\epsilon$ -structured strategy for  $n$  sequential CHSH games, then there is a single-qubit ideal strategy  $\tilde{\mathcal{S}}$  that  $\kappa n^\kappa \epsilon^{1/\kappa}$ -simulates an isometric extension of  $\mathcal{S}$ .*

*Proof sketch.* As explained in the main text, let  $\tilde{\mathcal{E}}_j^D$  be the super-operator that replaces the device's measurement operators with the ideal operators promised by the CHSH rigidity theorem. Then  $\tilde{\mathcal{S}} = (\rho_1, \{\tilde{\mathcal{E}}_j^A\}, \{\tilde{\mathcal{E}}_j^B\})$ . If  $\Pr[\text{game } j \text{ is } \epsilon\text{-structured}] \geq 1 - \delta$ , then  $\|\mathcal{E}_j^{AB}(\rho_j) - \tilde{\mathcal{E}}_j^{AB}(\rho_j)\|_{\text{tr}} \leq 2\delta + O(\sqrt{\epsilon})$ . (This expression combines bounds on the probability of the bad event and the  $O(\sqrt{\epsilon})$  error from the good event.)

To show our goal, that  $\mathcal{E}_{1,n}^{AB}(\rho_1) \approx \tilde{\mathcal{E}}_{1,n}^{AB}(\rho_1)$  in trace distance, use a hybrid argument that works backwards from game  $n$  to game 1 fixing each game's measurement operators one at a time. The error introduced from fixing a game  $j$ , by moving from  $\mathcal{E}_j^{AB}(\rho_j)$  to  $\tilde{\mathcal{E}}_j^{AB}(\rho_j)$ , does not increase in later games because applying a super-operator cannot increase the trace distance. Mathematically, this hybrid argument is simply a triangle inequality using the

expansion

$$\mathcal{E}_{1,n}^{AB}(\rho_1) - \tilde{\mathcal{E}}_{1,n}^{AB}(\rho_1) = \sum_{j \in [n]} \tilde{\mathcal{E}}_{j+1,n}^{AB}(\mathcal{E}_j^{AB}(\rho_j) - \tilde{\mathcal{E}}_j^{AB}(\rho_j)).$$

Therefore  $\tilde{\mathcal{S}}$  weakly simulates  $\mathcal{S}$ ; the extension to simulation is more technical, as described in the main text.  $\square$

Next, we find a nearby strategy in which the qubits for successive games are in tensor product.

**Definition B.8** (Multi-qubit ideal strategy). *A strategy  $\mathcal{S}$  is a multi-qubit ideal strategy if there is a unitary isomorphism  $\mathcal{Y}^D : \mathcal{H}_D \xrightarrow{\cong} (\mathbf{C}^2)^{\otimes n} \otimes \mathcal{H}'_D$  under which for unitaries  $M_j^D(h_{j-1}^D) \in \mathcal{L}((\mathbf{C}^2)^{\otimes(n-j+1)} \otimes \mathcal{H}'_D)$  such that*

$$R_\alpha^D(h_{j-1}^D) = \mathcal{Y}^{D\dagger} M_1^{D\dagger} \dots (\mathbf{1}_{\mathbf{C}^2}^{\otimes(j-1)} \otimes M_j^D(h_{j-1}^D)^\dagger) (R_\alpha^D)_j (\mathbf{1}_{\mathbf{C}^2}^{\otimes(j-1)} \otimes M_j^D(h_{j-1}^D)) \dots M_1^D \mathcal{Y}^D. \quad (8)$$

*That is,  $\mathcal{S}$  is a single-qubit ideal strategy in which the qubits used in each game must lie in tensor product with the qubits from previous games.*

**Theorem B.9.** *There exists a constant  $\kappa$  such that if  $\tilde{\mathcal{S}}$  is an  $\epsilon$ -structured single-qubit ideal strategy for  $n$  sequential CHSH games, then there is a multi-qubit ideal strategy  $\hat{\mathcal{S}}$  that  $\kappa n^\kappa \epsilon^{1/\kappa}$ -simulates an isometric extension of  $\tilde{\mathcal{S}}$ .*

*Proof sketch.* The tensor-product structure is constructed beginning with a trivial transformation on  $\tilde{\mathcal{S}}$ : to each device, add  $n$  ancilla qubits each in state  $|0\rangle$ . Next, after a qubit has been measured, say as  $|\alpha_j\rangle$  in game  $j$ , swap it

with the  $j$ th ancilla qubit, then rotate this fresh qubit from  $|0\rangle$  to  $|\alpha_j\rangle$  and continue playing games  $j+1, \dots, n$ . This defines a unitary change of basis that places the outcomes for games 1 to  $j$  in the first  $j$  ancilla qubits, and leaves the state in the original Hilbert space unchanged. Since qubits are set aside after being measured, the qubits for later games are automatically in tensor product with those for earlier games; the resulting strategy  $\tilde{\mathcal{S}}$  is multi-qubit ideal. At the end of the  $n$  games, swap back the ancilla qubits and undo their rotations, using the transcript.

The key to showing that  $\tilde{\mathcal{S}}$  is close to  $\tilde{\mathcal{S}}$  is the fact that operations on one half of an EPR state can equivalently be performed on the other half, since for any  $2 \times 2$  matrix  $M$ ,  $(M \otimes I)(|00\rangle + |11\rangle) = (I \otimes M^T)(|00\rangle + |11\rangle)$ . This means that the outcome of an  $\epsilon$ -structured CHSH game would be nearly unchanged if Bob were hypothetically to perform Alice's measurement before his own. By moving Alice's measurement operators for games  $j+1$  to  $n$  over to Bob's side, we see that they cannot significantly affect the qubit  $|\alpha_j\rangle$  from game  $j$  on her side. Therefore, undoing the original change of basis restores the ancilla qubits nearly to their initial state  $|0^n\rangle$ , and  $\tilde{\mathcal{S}} \approx \tilde{\mathcal{S}}$ .

Formally, define a unitary super-operator  $\mathcal{V}_j$  that rotates the  $j$ th ancilla qubit to  $|\alpha_j\rangle$ , depending on Alice's local transcript  $h_j^A$ . Define a unitary super-operator  $\mathcal{T}_j$  to apply  $\mathcal{V}_j$  and swap the  $j$ th ancilla qubit with the qubit Alice uses in game  $j$  (depending on  $h_{j-1}^A$ ). Alice's multi-qubit ideal strategy is given by

$$\tilde{\mathcal{E}}_j^A = \mathcal{T}_{1,j-1}^{-1}(\mathbf{1}_{\mathbb{C}^{2^n}} \otimes \tilde{\mathcal{E}}_j^A) \mathcal{T}_{1,j-1} . \quad (9)$$

We aim to show that the strategy given by  $\rho_1$ ,  $\{\tilde{\mathcal{E}}_j^A\}$  and  $\{\tilde{\mathcal{E}}_j^B\}$  is close to  $\tilde{\mathcal{S}}$  up to the fixed isometry that prepends  $|0^n\rangle\langle 0^n|$  to the state, i.e., that  $|0^n\rangle\langle 0^n| \otimes \tilde{\mathcal{E}}_{1,n}^{AB}(\rho_1) \approx \tilde{\mathcal{E}}_{1,n}^A(|0^n\rangle\langle 0^n| \otimes \tilde{\mathcal{E}}_{1,n}^B(\rho_1))$ . Define a super-operator  $\tilde{\mathcal{F}}_j^{AB}$ , in which Alice's measurements are made on Bob's Hilbert space  $\mathcal{H}_B$ , on the qubit determined by Bob's local transcript  $h_{j-1}^B$ . Since most games are  $\epsilon$ -structured, by the CHSH rigidity theorem,  $\tilde{\mathcal{F}}_{j+1,k}^{AB}(\tilde{\rho}_{j+1}) \approx \tilde{\mathcal{E}}_{j+1,k}^{AB}(\tilde{\rho}_{j+1}) = \tilde{\rho}_{k+1}$  for any  $j \leq k$ . Since  $\tilde{\mathcal{F}}_{j+1,k}^{AB}$  acts on  $\mathcal{H}_B$ , it does not affect Alice's qubit  $|\alpha_j\rangle$  from game  $j$  at all, and so this qubit must stay near  $|\alpha_j\rangle$  in  $\tilde{\rho}_{k+1}$  as well, i.e., the trace of the reduced density matrix against the projection  $|\alpha_j\rangle\langle \alpha_j|$  stays close to one. As this holds for every  $j$ ,  $\mathcal{T}_{1,n}^{-1}$  indeed returns the ancillas almost to their initial state  $|0^n\rangle$ .

In more detail, let  $X_j$  be the operator that projects onto Alice's  $j$ th ancilla qubit and the qubit she uses in the  $j$ th game being  $|0\rangle \otimes |\alpha_j\rangle$ . By definition,  $\text{Tr}(X_j \tilde{\rho}_{j+1}) = 1$ . By the Gentle Measurement Lemma,<sup>40,41</sup> it suffices to show that  $\text{Tr}(X_j \tilde{\rho}_{k+1}) = \text{Tr} X_j \tilde{\mathcal{E}}_{j+1,k}^{AB}(\tilde{\rho}_{j+1}) \approx 1$ . This is not obvious; since the operators for games  $j+1$  to  $k$  do not act in tensor product, they can disturb the qubit measured in game  $j$ . However, since a super-operator on  $\mathcal{H}_B$  cannot affect the expectation of an operator supported on  $\mathcal{H}_A$ , we

find

$$\begin{aligned} \text{Tr}(X_j \tilde{\rho}_{k+1}) &= \text{Tr} X_j \tilde{\mathcal{E}}_{j+1,k}^{AB}(\tilde{\rho}_{j+1}) \\ &\approx \text{Tr} X_j \tilde{\mathcal{F}}_{j+1,k}^{AB}(\tilde{\rho}_{j+1}) \\ &= \text{Tr}(X_j \tilde{\rho}_{j+1}) \\ &= 1 . \end{aligned}$$

A symmetrical argument adjusts Bob's super-operators  $\{\tilde{\mathcal{E}}_j^B\}$  to  $\{\tilde{\mathcal{E}}_j^A\}$ , implying that  $\tilde{\mathcal{S}}$  weakly simulates  $\tilde{\mathcal{S}}$ .  $\square$

The last major step in the proof of [Theorem B.5](#) is to argue that the qubit locations cannot depend significantly on the local transcripts, and therefore simulate a multi-qubit ideal strategy with an ideal strategy.

**Theorem B.10.** *There exists a constant  $\kappa$  such that if  $\tilde{\mathcal{S}}$  is an  $\epsilon$ -structured multi-qubit ideal strategy for  $n$  sequential CHSH games, then there is an ideal strategy that  $\kappa n^\kappa \epsilon^{1/\kappa}$ -simulates  $\tilde{\mathcal{S}}$ .*

*Proof sketch.* Fix a transcript  $\hat{h}_n$ , chosen at random from the distribution of transcripts for  $\tilde{\mathcal{S}}$ . Define a strategy  $\hat{\mathcal{S}}$  to use the same initial state  $\bar{\rho}_1$  as  $\tilde{\mathcal{S}}$  and the qubits specified by  $\hat{h}_n$  in  $\tilde{\mathcal{S}}$ , i.e.,  $\hat{R}_\alpha^D(h_{j-1}^D) = \hat{R}_\alpha^D(\hat{h}_{j-1}^D)$  independent of  $h_{j-1}^D$ . Thus, as required in [Definition B.2](#),  $\hat{R}_\alpha^D(h_{j-1}^D) = \hat{\mathcal{I}}^{D\dagger}(R_\alpha^D)_j \hat{\mathcal{I}}^D$  for

$$\hat{\mathcal{I}}^D = (\mathbf{1}_{(\mathbb{C}^2)^{\otimes (n-1)}} \otimes \bar{M}_n^D(h_{n-1}^D)) \dots \bar{M}_1^D \bar{\mathcal{Y}}^D .$$

We argue that  $\hat{\mathcal{S}}$  closely approximates  $\tilde{\mathcal{S}}$ , provided that  $\hat{h}_n$  satisfies: for every  $j$ , conditioned on the partial transcript  $\hat{h}_{j-1}$ , (a) game  $j$  is  $\epsilon$ -structured, and (b) there is a high probability that every subsequent game is  $\epsilon$ -structured. By Markov inequalities, most transcripts satisfy these conditions.

We connect  $\tilde{\mathcal{S}}$  to  $\hat{\mathcal{S}}$  by an argument that one game at a time switches play to locate qubits according to  $\hat{h}_n$ . The intermediate steps relate strategies in which the devices locate their qubits using a hybrid  $(\hat{h}_j, h_{j+1,n})$  of  $\hat{h}_n$  and the actual transcript  $h_n$ .

Consider a partial transcript  $h_j$  that differs from  $\hat{h}_j$  only in the  $j$ th game, say on Alice's side. By (a) and the CHSH rigidity theorem, Alice's  $j$ th qubit is collapsed and nearly in tensor product with the rest of the state. Therefore, there exists a unitary  $V_j^A$  acting on this qubit such that

$$\bar{\rho}_{j+1}(h_j) \approx V_j^A \bar{\rho}_{j+1}(\hat{h}_j) V_j^{A\dagger} , \quad (10)$$

up to error  $O(\sqrt{\epsilon})$ . Since applying a super-operator cannot increase trace distance and on Bob's side  $\hat{h}_j^B = h_j^B$ , therefore

$$\tilde{\mathcal{F}}_{j+1,n}^{AB|h_j^B}(\bar{\rho}_{j+1}(h_j)) \approx V_j^A \tilde{\mathcal{F}}_{j+1,n}^{AB|\hat{h}_j^B}(\bar{\rho}_{j+1}(\hat{h}_j)) V_j^{A\dagger} .$$

Here,  $\tilde{\mathcal{F}}_{j+1,n}^{AB|h_j^B}$  is the same super-operator used in the multi-qubit ideal strategy simulation step—that plays Alice's

games on Bob's qubits—except conditioned on the local transcript  $h_j^B$ . By condition (b), these super-operators can be pulled back to Alice's side, to give

$$\bar{\mathcal{E}}_{j+1,n}^{AB|h_j}(\bar{\rho}_{j+1}(h_j)) \approx V_j^A \bar{\mathcal{E}}_{j+1,n}^{AB|\hat{h}_j}(\bar{\rho}_{j+1}(\hat{h}_j)) V_j^{A\dagger}.$$

Note that this approximation does not follow immediately from Eq. (10), because Alice's super-operators conditioned on  $h_j^A$  can be very different from her super-operators conditioned on  $\hat{h}_j^A$ .

By fixing the coordinates one at a time in this way, we find that for a typical transcript  $h_n$ ,  $\bar{\rho}_{n+1}(h_n) \approx V_{1,n}^{AB} \bar{\rho}_{n+1}(\hat{h}_n) V_{1,n}^{AB\dagger}$ , and we conclude that  $\bar{\mathcal{E}}_{1,n}^{AB}(\rho_1) \approx \hat{\mathcal{E}}_{1,n}^{AB}(\rho_1)$ .

Since  $\hat{\mathcal{E}}_{1,n}^{AB}$  measures qubits in tensor product with each other, by using the CHSH rigidity theorem one last time, it is not difficult to show that  $\hat{\mathcal{E}}_{1,n}^{AB}(\rho_1) \approx \hat{\mathcal{E}}_{1,n}^{AB}(\hat{\rho}_1)$ , where  $\hat{\rho}_1$  has  $n$  EPR pairs in the qubit positions determined by  $\hat{h}_n$ . Thus  $\hat{\mathcal{S}}$  is weakly simulated by the ideal strategy given by  $(\hat{\rho}_1, \{\hat{\mathcal{E}}_j^A\}, \{\hat{\mathcal{E}}_j^A\})$ .  $\square$

These three steps chain together to yield [Theorem B.5](#), via:

**Lemma B.11.** *Let  $\mathcal{S}$  be a  $(\delta, \epsilon)$ -structured strategy for  $n$  sequential CHSH games. If  $\tilde{\mathcal{S}}$  is a strategy that weakly  $\eta$ -simulates  $\mathcal{S}$ , then  $\tilde{\mathcal{S}}$  is  $(\delta + 2\sqrt{\eta}, \epsilon + 16\sqrt{\eta})$ -structured.*

This lemma follows immediately from the definitions.

## C Verified quantum dynamics proof sketches

### C.1 XZ-determined states

In the CHSH game, each device has two measurement settings, that in the ideal strategy may be identified with Pauli  $\sigma_x$  and  $\sigma_z$  operators. For carrying out tomography, however, it is generally necessary to be able to measure in the  $\sigma_y$  basis as well. It is possible to extend the CHSH game to one in which the ideal strategy also uses  $\sigma_y$  operators on a shared EPR state. However, this extended game will not satisfy the rigidity property. The problem is that a device that consistently measures using  $-\sigma_y$  will give indistinguishable statistics from one that uses  $+\sigma_y$ . (Switching the sign of  $\sigma_y$  is equivalent to taking an entry-wise complex conjugate in the computational basis.) It is impossible to fix the sign of the  $\sigma_y$  operator.

We therefore instead argue that for certain states, reliable tomography can be accomplished without needing to measure in the  $\sigma_y$  basis. This observation is due to McKague<sup>36</sup> and was suggested earlier by Magniez et al.<sup>42</sup> McKague shows that for  $|\varphi\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$ , an EPR pair, the states  $(I \otimes U)|\varphi\rangle$ , for any single-qubit real unitary  $U$ , and  $\text{CNOT}_{24}|\varphi\rangle_{12} \otimes |\varphi\rangle_{34}$ , as well as finite tensor products of these states, are exactly determined by their traces against tensor products of  $I$ ,  $\sigma_x$  and  $\sigma_z$  operators.

That is, they are determined by the expectations of observables that can be estimated using measurements in the  $\sigma_x$  and  $\sigma_z$  bases. We call such states “XZ-determined.” For our applications, we will need to show a larger class of states to be XZ-determined. However, characterizing the full set of XZ-determined states remains an open problem.

**Definition C.1.** *For a Hilbert space  $\mathcal{H}$ , a set of operators  $S \subseteq \mathcal{L}(\mathcal{H})$  and  $d > 0$ , a state  $\tau \in \mathcal{L}(\mathcal{H})$  is determined by  $S$  with exponent  $d$  if there exists  $c > 0$  such that for all  $\epsilon \geq 0$  and any state  $\rho \in \mathcal{L}(\mathcal{H})$ ,*

$$\max_{P \in S} |\text{Tr } P(\rho - \tau)| \leq \epsilon \implies \|\rho - \tau\|_{\text{tr}} \leq c\epsilon^d. \quad (11)$$

*For  $\mathcal{H} = (\mathbf{C}^2)^{\otimes n}$ , a state  $\tau$  is XZ-determined if it is determined with some exponent  $d > 0$  by the Pauli operators  $\{I, \sigma_x, \sigma_z\}^{\otimes n}$ .*

Robustness is important for applications, but previous work has considered only  $\epsilon = 0$ . By Łojasiewicz's inequality [43, Prop. 2.3.11] in algebraic geometry, robustness follows from the  $\epsilon = 0$  case:

**Lemma C.2.** *For a finite-dimensional Hilbert space  $\mathcal{H}$ , a state  $\tau \in \mathcal{L}(\mathcal{H})$  is determined by a finite set  $S \subset \mathcal{L}(\mathcal{H})$  if and only if for any state  $\rho \in \mathcal{L}(\mathcal{H})$ , the implication of Eq. (11) holds at  $\epsilon = 0$ .*

Recall that a *stabilizer state* is an  $n$ -qubit pure state  $|\psi\rangle$  for which there exists a set of  $2^n$  distinct and pairwise commuting operators  $S \subset \{\pm P : P \in \{I, \sigma_x, \sigma_y, \sigma_z\}^{\otimes n}\}$ , the *stabilizer group*, such that  $P|\psi\rangle = |\psi\rangle$  for all  $P \in S$ .<sup>44</sup> Any set of  $n$  operators that generate the stabilizer group  $S$  are called *stabilizer generators* for  $|\psi\rangle$ . We prove:

**Theorem C.3.** *If  $|\psi\rangle \in (\mathbf{C}^2)^{\otimes n}$  is a stabilizer state that has a set of stabilizer generators in  $\{I, \sigma_x, \sigma_z\}^{\otimes n}$ , and if  $U$  is the tensor product of any  $n$  single-qubit real unitaries, then  $U|\psi\rangle$  is XZ-determined.*

In particular, the resource states needed in our verified quantum dynamics protocol,  $|0\rangle \otimes (I \otimes H)|\varphi\rangle \otimes (I \otimes G)|\varphi\rangle \otimes \text{CNOT}_{2,4}(|\varphi\rangle \otimes |\varphi\rangle)$ , are XZ-determined, as are the same states with  $G$  applied transversally. This latter consideration is important because Alice's ideal measurement bases in the CHSH game are rotated by  $\pi/4$  from  $\sigma_x$  and  $\sigma_z$ .

The proof of [Theorem C.3](#) begins by showing easily that  $|0\rangle$  is determined by  $\{\sigma_z\}$ . Then apply three closure properties. First, closure under tensor product for pure states implies that  $|0\rangle^{\otimes n}$  is determined by the  $n$  operators  $\sigma_z \otimes I^{\otimes(n-1)}, \dots, I^{\otimes(n-1)} \otimes \sigma_z$ . Second, observe by manipulating implication (11) that if  $\tau$  is a state determined by  $S$ , then for any unitary  $U$ ,  $U\tau U^\dagger$  is determined with the same exponent by  $\{UPU^\dagger : P \in S\}$ . Since an arbitrary stabilizer state can be generated by applying Clifford operators to  $|0\rangle^{\otimes n}$ , this implies that a stabilizer state is determined by any of its sets of stabilizer generators. Furthermore, if  $\tau$  is determined by  $S = \{P_1, \dots, P_s\}$ , then for any invertible

$s \times s$  matrix  $V$ ,  $\tau$  is determined with the same exponent by  $\{\sum_{j=1}^s V_{ij} P_j : i = 1, \dots, s\}$ . This implies, e.g., that any  $XZ$ -determined state is also determined by the operators  $\{I, \frac{1}{\sqrt{2}}(\sigma_z \pm \sigma_x)\}^{\otimes n}$ , since the  $\{I, \sigma_x, \sigma_z\}^{\otimes n}$  coefficients are functions of the  $\{I, \frac{1}{\sqrt{2}}(\sigma_z \pm \sigma_x)\}^{\otimes n}$  coefficients.

## C.2 State tomography

The state tomography protocol begins with  $(K - 1)m$  CHSH games, where  $K$  is chosen randomly. The multi-game CHSH rigidity theorem implies that at the beginning of the  $K$ th block of  $m$ , with high probability Alice and Bob share a state that is close to  $m$  shared EPR states, possibly in tensor product with an additional state, and that their separate measurement strategies for the next  $m$  games are close to the ideal strategy that uses one EPR state at a time. For the analysis of the state tomography protocol, we may therefore assume that Alice's strategy is exactly ideal and restrict consideration to these  $m$  EPR states.

Bob, on the other hand, does not play more CHSH games, but instead is given by Eve a random permutation of the  $m$  indices, and is asked to permute his qubits and prepare many copies of a particular resource state. (In the main text, this state is specified as  $|\psi\rangle = |0\rangle \otimes (I \otimes H)|\varphi\rangle \otimes (I \otimes G)|\varphi\rangle \otimes \text{CNOT}_{2,4}(|\varphi\rangle \otimes |\varphi\rangle)$ . The  $|0\rangle$  portion is for preparing the initial states in a computation and implementing the final measurements, and the other subsystems are for teleporting into each of the gates in a universal gate set, e.g.,  $(I \otimes G)|\varphi\rangle$  for teleporting into  $G$ . However, after teleporting into  $G$ , an  $H$  correction may or may not be required. To maintain the blindness property of the protocol, i.e., to avoid leaking any information about the computation to the separate devices, it is of technical use to have available the resource state  $|\psi\rangle \otimes |\varphi\rangle$ , where the extra EPR state is used for teleporting into the identity gate when an  $H$  correction is not needed.)

Note that Bob's reduced density matrix is maximally mixed, so the probability that he can measure the correct 11-qubit resource state is only  $1/2^{11}$ . However, since the states  $\{(P \otimes I)|\varphi\rangle : P \in \{I, \sigma_x, \sigma_y, \sigma_z\}\}$  form an orthonormal basis, so too do the states

$$P^{(0)}|0\rangle \otimes (P^{(1)} \otimes I)|\varphi\rangle \otimes (P^{(2)} \otimes H)|\varphi\rangle \otimes (P^{(3)} \otimes G)|\varphi\rangle \\ \otimes (P_1^{(4)} \otimes P_3^{(5)} \otimes \text{CNOT}_{2,4})(|\varphi\rangle_{12} \otimes |\varphi\rangle_{34}), \quad (12)$$

where  $P^{(0)} \in \{I, \sigma_x\}$  and the other  $P^{(j)}$  vary over  $\{I, \sigma_x, \sigma_y, \sigma_z\}$ . Any of the states in Eq. (12) are equally useful resources for computation by teleportation, as Eve can adjust for the  $P^{(j)}$  operators in her classical Pauli frame.<sup>45</sup>

Therefore define an ideal state tomography protocol as one in which Alice and Bob's initial state consists of  $m$  shared EPR states, possibly in tensor product with an additional state; Alice plays honestly  $m$  CHSH games, directed

by Eve; and Eve sends Bob a random  $m$ -item permutation and requests that he return the results of measuring consecutive 11-qubit blocks of permuted qubits in the basis of Eq. (12). Eve rejects if the tomography statistics returned by Alice are inconsistent with Bob's reported outcomes. More precisely, she rejects if the fraction of times Bob reports any particular state differs from  $1/2^{11}$  by more than  $\sqrt{(\log m)/m}$ , i.e., about  $\sqrt{\log m}$  standard deviations, or if for any state any of its  $\{I, X, Z\}^{\otimes 11}$  Pauli coefficients differ from the corresponding observed estimates by more than  $\sqrt{(\log m)/m}$ . We show:

**Theorem C.4.** *In an ideal state tomography protocol, if Alice plays honestly, then:*

**Completeness:** *If Bob plays honestly, then Eve accepts with high probability.*

**Soundness:** *If Eve accepts with high probability, then there is a high probability that, after Bob and before Alice's play, for most of the consecutive 11-qubit subsystems, Alice's reduced state on the subsystem is close to the state in Eq. (12) that Bob reported to Eve.*

In these statements, "with high probability" means with probability inverse polynomially close to one, i.e., at least  $1 - 1/m^c$ , where the exact exponents are adjustable.

The completeness property of the protocol is a trivial application of Hoeffding's inequality; the probability of straying by more than  $\sqrt{\log m}$  standard deviations is at most  $\exp(-\Omega(\log m)) = m^{-\Omega(1)}$ . The soundness property is also mostly a straightforward tomography argument, using that the states in Eq. (12) are all  $XZ$ -determined. The main technical complication is that the states of Alice's 11-qubit blocks need not be in tensor product. Therefore, her measurement results on different blocks need not be independent. They can be controlled with a suitable martingale.

First fix a permutation  $\sigma \in S_m$  and a string  $x \in \{0, 1\}^{11m/11}$  such that, conditioned on Eve sending Bob  $\sigma$  and receiving back  $x$ , Eve accepts with high probability. The remaining randomness consists of Alice's measurement bases and results. Since Alice's measurements commute, we may assume that she measures her qubits in the permuted order, without changing the measurement statistics.

For  $j = 1, \dots, m/11$ , let  $\rho_j$  be Alice's initial reduced state on her  $j$ th block of 11 qubits. Our goal is to control most of the states  $\rho_j$ . Let  $\sigma_j$  be the state of the same qubits just before she begins to measure them. The state  $\sigma_j$  is a random variable, but is a deterministic function of the transcript  $h_{11(j-1)}^A$  of the earlier CHSH games with Alice. Conditioned on  $h_{11(j-1)}^A$ , Alice's measurement results for games  $11(j-1) + 1, \dots, 11j$  are distributed according to the Pauli coordinates of  $\sigma_j$ .

For each  $b \in \{0, 1\}^{11}$ , let  $\pi_b$  be the corresponding state of Eq. (12), and let  $\tau_b$  be the average of those  $\sigma_j$  for which Bob reported  $x_j = b$ . Using a martingale argument, we



can establish that with high probability, for all  $b$ ,  $\pi_b$  and  $\tau_b$  have similar  $\{I, \sigma_x, \sigma_z\}^{\otimes 11}$  Pauli coordinates. Since  $\pi_b$  is  $XZ$ -determined, this implies that all Pauli coordinates of  $\pi_b$  and  $\tau_b$  are close. Since  $\pi_b$  is a pure state, i.e., an extremal quantum state, a Markov inequality implies that for most  $j$  with  $x_j = b$ ,  $\pi_b$  and  $\sigma_j$  have close Pauli coordinates. Finally, average back over the transcripts to get that for most  $j$  with  $x_j = b$ ,  $\pi_b$  is close to  $\rho_j = \sum_{h_{11(j-1)}^A} \Pr[h_{11(j-1)}^A] \sigma_j(h_{11(j-1)}^A)$ . This is the claim.

**Theorem C.4** says that Bob’s measurement usually prepares the correct state on Alice’s side—but it says nothing about the distribution of his measurement results. Since Bob’s half of the shared EPR states is maximally mixed, his measurement outcomes are in fact distributed nearly uniformly on most subsystems. Thus the effect of Bob’s actual super-operator is close to that of the ideal super-operator, if we trace out everything except for a random subset of subsystems on Alice’s side.

### C.3 Process tomography

Computation by teleportation uses adaptively chosen two-qubit Bell measurements on prepared resource states. The state tomography protocol gives Eve a way of ensuring that Alice’s initial  $m$ -qubit state consists of the desired resource states. The Bell states  $\{(P \otimes I)|\varphi\rangle : P \in \{I, \sigma_x, \sigma_y, \sigma_z\}\}$ , eigenstates of the Bell measurement, are themselves  $XZ$ -determined states. A symmetrical protocol, with the roles of Alice and Bob switched, could thus be used to prepare these states in Bob’s initial  $m$ -qubit state, before his play begins. However, for verified quantum computation we need a stronger conclusion, that *after* Bob’s measurements collapse Alice’s initial state into copies of the resource states of Eq. (12), Alice’s operations indeed implement Bell measurements. This is a stronger conclusion because, as explained earlier, it is possible that dishonest operations on Alice’s side—far from Bell measurements—could still collapse Bob’s initial state to many copies of Bell states.

The process tomography protocol, to ensure Alice’s honesty, begins in the same way as state tomography, with  $(K-1)m$  CHSH games refereed between both players. The multi-game rigidity theorem again allows us to assume that the devices then share  $m$  EPR states and would play the  $K$ th block of games honestly, according to the ideal CHSH game strategy. However, Eve only continues playing CHSH games with Bob. To Alice, Eve repeatedly sends random pairs of indices, drawn uniformly without replacement from  $\{1, \dots, m\}$ , requesting the results of Bell measurements on the specified qubits. Eve does not try to draw any statistical conclusions from the reported results, but instead rejects if Alice’s reported result for any pair of qubits is inconsistent with Bob’s reported outcomes. For example, if Alice reports measuring the Bell state  $|\varphi\rangle$  and Bob is asked to measure the same two qubits both in the  $\sigma_z$  basis, then the parity of Bob’s measurement results should be 0. This

is because  $\sigma_z \otimes \sigma_z$  is a stabilizer for  $|\varphi\rangle$ :  $(\sigma_z \otimes \sigma_z)|\varphi\rangle = |\varphi\rangle$ . The four Bell states are all stabilized by  $\pm\sigma_z \otimes \sigma_z$  and  $\pm\sigma_x \otimes \sigma_x$ , with four different sign choices. If Bob is asked to measure one of the two qubits in the  $\sigma_x$  basis and one in the  $\sigma_z$  basis, then his results for that pair do not give a check on Alice.

From Alice’s perspective, the interaction is indistinguishable from the questions she would be asked in an actual computation; she is entirely unaware that Bob is playing CHSH games instead of preparing resource states.

We show:

**Theorem C.5.** *In the process tomography protocol, if after the first  $(K-1)m$  CHSH games, Bob plays honestly on his halves of  $m$  shared EPR states, then:*

**Completeness:** *If Alice plays honestly, Eve accepts with probability one.*

**Soundness:** *If Eve accepts with high probability, then the result of Alice’s super-operator applied to the initial state is close to that of applying the ideal Bell measurements super-operator to the initial state.*

Again, the completeness statement is trivial. The soundness statement is not difficult. Fix a permutation of the qubits for which Eve accepts with high probability. We prove soundness for the protocol in which Alice is given the full permutation at the beginning instead of only two indices at a time; this can only give her more opportunities to cheat.

Let  $\rho_1$  be the initial state, consisting of  $m$  EPR states. Without loss of generality, Alice’s strategy consists of measuring a complete set of  $2^m$  orthogonal projections, and returning the outcome. For  $j = 1, \dots, m/2$ , let  $\mathcal{G}_j^A$  be Alice’s super-operator that implements the  $j$ th alleged Bell measurement, by the appropriate marginal projective measurement. For  $j \leq k$ , let  $\mathcal{G}_{j,k}^A = \mathcal{G}_k^A \dots \mathcal{G}_{j+1}^A \mathcal{G}_j^A$ . Alice’s full strategy is implemented by  $\mathcal{G}_{1,m/2}^A$ . Let  $\hat{\mathcal{G}}_j^A$  be the ideal super-operator that actually carries out a Bell measurement on the  $j$ th specified pair of qubits, and  $\hat{\mathcal{G}}_{j,k}^A = \hat{\mathcal{G}}_k^A \dots \hat{\mathcal{G}}_j^A$ . Our goal is to show that in trace distance

$$\mathcal{G}_{1,m/2}^A(\rho_1) \approx \hat{\mathcal{G}}_{1,m/2}^A(\rho_1) . \quad (13)$$

Since Eve accepts all the tests with high probability, it must be that for every  $j$ , Eve’s  $j$ th Bell measurement test passes with high probability. By the Gentle Measurement Lemma, this implies that  $\mathcal{G}_j^A(\rho_1) \approx \hat{\mathcal{G}}_j^A(\rho_1)$ . For  $j = 1$ , this gives the first step:

$$\mathcal{G}_{1,m/2}^A(\rho_1) \approx \mathcal{G}_{2,m/2}^A \hat{\mathcal{G}}_1^A(\rho_1) .$$

We cannot immediately apply  $\mathcal{G}_2^A(\rho_1) \approx \hat{\mathcal{G}}_2^A(\rho_1)$  to continue, because although the  $\mathcal{G}_j^A$  super-operators commute with each other,  $\mathcal{G}_2^A$  might not commute with  $\hat{\mathcal{G}}_1^A$ .

An easy trick gets around the problem. Define  $\hat{\mathcal{F}}_j^B$  to implement a Bell measurement on the  $j$ th specified pair of qubits on Bob's side. Let  $\hat{\mathcal{F}}_{j,k}^B = \hat{\mathcal{F}}_k^B \cdots \hat{\mathcal{F}}_j^B$ . Then  $\hat{\mathcal{G}}_j^A(\hat{\rho}_1) = \hat{\mathcal{F}}_j^B(\hat{\rho}_1)$ . Super-operators acting on Bob's Hilbert space automatically commute with those acting on Alice's space, so we can continue the above derivation:

$$\begin{aligned} \mathcal{G}_{1,m/2}^A(\hat{\rho}_1) &\approx \mathcal{G}_{2,m/2}^A \hat{\mathcal{G}}_1^A(\hat{\rho}_1) = \hat{\mathcal{F}}_1^B \mathcal{G}_{2,m/2}^A(\hat{\rho}_1) \\ &\approx \hat{\mathcal{F}}_1^B \mathcal{G}_{3,m/2}^A \hat{\mathcal{G}}_2^A(\hat{\rho}_1) = \hat{\mathcal{F}}_{1,2}^B \mathcal{G}_{3,m/2}^A(\hat{\rho}_1) \approx \cdots \\ &\approx \hat{\mathcal{F}}_{1,m/2}^B(\hat{\rho}_1) = \hat{\mathcal{G}}_{1,m/2}^A(\hat{\rho}_1) . \end{aligned}$$

The total approximation error is linear in  $m$ .

#### C.4 QMIP = MIP\*

The way in which protocols for CHSH games, state and process tomography, and computation are combined to give verified quantum dynamics is sketched in the main text. We also describe there the main remaining technical obstacle, the issue of Eve choosing her questions adaptively. Formally, let  $\rho$  be the initial state, and let  $\mathcal{B}$  be the super-operator describing Eve's interactions with Bob in state tomography. Roughly, state tomography implies that the states Bob prepares on Alice's side are correct up to a small error in trace distance, or

$$\text{Tr}_B \mathcal{B}(\rho) \approx \text{Tr}_B \hat{\mathcal{B}}(\hat{\rho}) , \quad (14)$$

where  $\hat{\mathcal{B}}$  is the ideal super-operator and  $\hat{\rho}$  is an ideal initial state consisting of perfect EPR states. Similarly, let  $\mathcal{A}$  be the super-operator describing Eve's interactions with Alice in a process tomography protocol on Alice's operations; we have

$$\mathcal{A}(\rho) \approx \hat{\mathcal{A}}(\rho) . \quad (15)$$

Computation by teleportation can be implemented either by choosing Bob's state preparation questions non-adaptively and Alice's process questions adaptively, or vice versa. We show that these are exactly equivalent regardless of the devices' strategies, i.e.,

$$\mathcal{A}_{\text{ad}} \mathcal{B} = \mathcal{B}_{\text{ad}} \mathcal{A} , \quad (16)$$

where  $\mathcal{A}_{\text{ad}}$  and  $\mathcal{B}_{\text{ad}}$  are the same as  $\mathcal{A}$  and  $\mathcal{B}$ , respectively, except with Eve choosing her questions adaptively based on the previous messages. Combining these steps, we therefore obtain

$$\begin{aligned} \text{Tr}_B \mathcal{B}_{\text{ad}} \mathcal{A}(\rho) &\approx \text{Tr}_B \mathcal{B}_{\text{ad}} \hat{\mathcal{A}}(\rho) \\ &= \hat{\mathcal{A}}_{\text{ad}} \text{Tr}_B \mathcal{B}(\rho) \\ &\approx \hat{\mathcal{A}}_{\text{ad}} \text{Tr}_B \hat{\mathcal{B}}_{\text{ad}}(\hat{\rho}) , \end{aligned}$$

and thus the actual computation by teleportation protocol leaves on Alice's side nearly the ideal output.

In this supplement, we would like to highlight two points of the proof that QMIP = MIP\*: the conversion into a

three-round protocol, and the addition of two new provers instead of reusing provers.

QMIP is the class of languages decidable by a polynomial-time quantum verifier exchanging polynomially many quantum messages with a polynomial number of quantum provers, who have unbounded computational power and share entanglement but cannot communicate among themselves. Kempe et al. have shown that any QMIP protocol can be converted into a three-turn protocol in which the provers send a quantum message to the verifier, the verifier broadcasts the result of a random coin flip, the provers each send a second quantum message, and then the verifier applies an efficient measurement to decide whether to accept or reject.<sup>38</sup> Beginning with this protocol transformation, our proof adds two new provers, Alice and Bob. The classical verifier, Eve, teleports both rounds of messages from the original  $k$  provers to Alice, and then directs Alice and Bob to run the original verifier's quantum circuit.

A natural question is whether it is necessary to add two new provers or if two of the provers already present can be used for implementing verified quantum dynamics. We conjecture that adding new provers is not necessary. Broadbent et al., for example, have suggested that a  $k$ -prover QMIP protocol can be converted to a  $k$ -prover MIP\* protocol deciding the same language, and the scheme that they present indeed reuses the first two provers to simulate the verifier's quantum computations.<sup>32</sup> However, the analysis of this scheme does not consider all ways in which provers can play dishonestly.

In fact, any scheme with a structure along the lines presented either by us or by Broadbent et al. will be unsound—if it does not first convert to a three-turn protocol or use more sophisticated tricks. Here is a general counterexample. Begin with an arbitrary QMIP protocol  $\mathcal{P}$ , deciding a language  $L$ . Modify the protocol by adding one new round at the end. In this last round, the provers can each send classical messages to the verifier, which the verifier simply broadcasts back to all of the provers. The effect of this final interaction is to allow the provers to communicate with each other. (Since they share entanglement, they can use quantum teleportation to communicate quantum information, if desired.) The modified protocol  $\mathcal{P}'$  decides the same language  $L$ , with the exact same completeness and soundness parameters, though; communicating in the last step does not help the provers cheat.

Convert the modified protocol, somehow, into an MIP\* protocol  $\mathcal{P}''$  that uses the first two provers to simulate the verifier's quantum computations. In particular, they simulate the final acceptance predicate, by some procedure that has traps to detect cheating—in our scheme, CHSH games or state or process tomography. The problem is that the last round of messages in  $\mathcal{P}'$  is useless, but in  $\mathcal{P}''$  the intercommunication allows the provers to reveal to each other the traps set by the verifier. This allows them to avoid the traps and cheat freely.  $\mathcal{P}''$  is unsound.

Converting to a three-turn protocol at the start deflects this general attack, as does using two fresh provers to run the verifier’s quantum computation. Speculatively, another way to avoid the attack might be to refresh the verifier’s secrets—resetting any traps and re-hiding any quantum information—before revealing any message to a prover. Any messages can carry information that affect the security of the converted MIP\* protocol differently from the original QMIP protocol.

## References

- [1] Bennett, C. H. & Brassard, G. Quantum cryptography: Public key distribution and coin tossing. In *Proc. IEEE Int. Conf. on Computers, Systems and Signal Processing*, 175 (1984).
- [2] Lo, H.-K. & Chau, H. F. Unconditional security of quantum key distribution over arbitrarily long distances. *Science* **283**, 2050–2056 (1999). arXiv:quant-ph/9803006.
- [3] Shor, P. W. & Preskill, J. Simple proof of security of the BB84 quantum key distribution protocol. *Phys. Rev. Lett.* **85**, 441–444 (2000). arXiv:quant-ph/0003004.
- [4] Zhao, Y., Fung, C.-H. F., Qi, B., Chen, C. & Lo, H.-K. Quantum hacking: experimental demonstration of time-shift attack against practical quantum key distribution systems. *Phys. Rev. A* **78**, 042333 (2008). arXiv:0704.3253 [quant-ph].
- [5] Lydersen, L. *et al.* Hacking commercial quantum cryptography systems by tailored bright illumination. *Nature Photonics* **4**, 686–689 (2010). arXiv:1008.4593 [quant-ph].
- [6] Gerhardt, I. *et al.* Full-field implementation of a perfect eavesdropper on a quantum cryptography system. *Nat. Commun.* **2**, 349 (2011). arXiv:1011.0105 [quant-ph].
- [7] Mayers, D. & Yao, A. Quantum cryptography with imperfect apparatus. In *Proc. 39th IEEE FOCS*, 503–509 (1998). arXiv:quant-ph/9809039.
- [8] Ekert, A. K. Quantum cryptography based on Bell’s theorem. *Phys. Rev. Lett.* **67**, 661–663 (1991).
- [9] Barrett, J., Hardy, L. & Kent, A. No signalling and quantum key distribution. *Phys. Rev. Lett.* **95**, 010503 (2005). arXiv:quant-ph/0405101.
- [10] Masanes, L., Renner, R., Christandl, M., Winter, A. & Barrett, J. Unconditional security of key distribution from causality constraints (2006). arXiv:quant-ph/0606049.
- [11] Acín, A., Massar, S. & Pironio, S. Efficient quantum key distribution secure against no-signalling eavesdroppers. *New J. Phys.* **8**, 126 (2006). arXiv:quant-ph/0605246.
- [12] Masanes, L. Universally composable privacy amplification from causality constraints. *Phys. Rev. Lett.* **102**, 140501 (2009). arXiv:0807.2158 [quant-ph].
- [13] Hänggi, E., Renner, R. & Wolf, S. Efficient device-independent quantum key distribution. In *Proc. EURO-CRYPT*, 216–234 (2010). arXiv:0911.4171 [quant-ph].
- [14] Acín, A. *et al.* Device-independent security of quantum cryptography against collective attacks. *Phys. Rev. Lett.* **98**, 230501 (2007). arXiv:quant-ph/0702152.
- [15] Pironio, S. *et al.* Device-independent quantum key distribution secure against collective attacks. *New J. Phys.* **11**, 045021 (2009). arXiv:0903.4460 [quant-ph].
- [16] McKague, M. Device independent quantum key distribution secure against coherent attacks with memoryless measurement devices. *New J. Phys.* **11**, 103037 (2009). arXiv:0908.0503 [quant-ph].
- [17] Hänggi, E. & Renner, R. Device-independent quantum key distribution with commuting measurements (2010). arXiv:1009.1833 [quant-ph].
- [18] Masanes, L., Pironio, S. & Acín, A. Secure device-independent quantum key distribution with causally independent measurement devices. *Nat. Commun.* **2**, 238 (2011). arXiv:1009.1567 [quant-ph].
- [19] Aharonov, D., Ben-Or, M. & Eban, E. Interactive proofs for quantum computations. In *Proc. Innovations in Computer Science (ICS)*, 453–469 (2010). arXiv:0810.5375 [quant-ph].
- [20] Broadbent, A., Fitzsimons, J. F. & Kashefi, E. Universal blind quantum computation. In *Proc. 50th IEEE FOCS*, 517–526 (2009). arXiv:0807.4154 [quant-ph].
- [21] Aharonov, D. & Vazirani, U. Is quantum mechanics falsifiable? A computational perspective on the foundations of quantum mechanics (2012). arXiv:1206.3686 [quant-ph].
- [22] Tsirelson, B. S. Quantum generalizations of Bell’s inequality. *Lett. Math. Phys.* **4**, 93–100 (1980).
- [23] Bell, J. S. On the Einstein-Podolsky-Rosen paradox. *Physics* **1**, 195–200 (1964).
- [24] Clauser, J. F., Horne, M. A., Shimony, A. & Holt, R. A. Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett.* **23**, 880–884 (1969).
- [25] Braunstein, S. L., Mann, A. & Revzen, M. Maximal violation of Bell inequalities for mixed states. *Phys. Rev. Lett.* **68**, 3259–3261 (1992).
- [26] Popescu, S. & Rohrlich, D. Which states violate Bell’s inequality maximally? *Phys. Lett. A* **169**, 411–414 (1992).
- [27] Bardyn, C.-E., Liew, T. C. H., Massar, S., McKague, M. & Scarani, V. Device independent state estimation based on Bell’s inequalities. *Phys. Rev. A* **80**, 062327 (2009). arXiv:0907.2170 [quant-ph].
- [28] McKague, M., Yang, T. H. & Scarani, V. Robust self-testing of the singlet (2012). arXiv:1203.2976 [quant-ph].
- [29] Miller, C. & Shi, Y. Robust self-testing quantum states and binary nonlocal XOR games (2012). arXiv:1207.1819 [quant-ph].
- [30] Fitzsimons, J. F. & Kashefi, E. Unconditionally verifiable blind computation (2012). arXiv:1203.5217 [quant-ph].
- [31] Barz, S. *et al.* Demonstration of blind quantum computing. *Science* **335**, 303–308 (2012). arXiv:1110.1381 [quant-ph].
- [32] Broadbent, A., Fitzsimons, J. F. & Kashefi, E. QMIP = MIP\* (2010). arXiv:1004.1130 [quant-ph].
- [33] Jordan, C. Essai sur la géométrie à  $n$  dimensions. *Bulletin de la S. M. F.* **3**, 103–174 (1875).
- [34] Reichardt, B. W., Unger, F. & Vazirani, U. A classical leash for a quantum system: Command of quantum systems via rigidity of CHSH games (2012). In preparation.
- [35] Gottesman, D. & Chuang, I. L. Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Nature* **402**, 390–393 (1999). quant-ph/9908010.
- [36] McKague, M. *Quantum information processing with adversarial devices*. Ph.D. thesis, University of Waterloo (2010). arXiv:1006.2352 [quant-ph].
- [37] Gottesman, D. *Stabilizer codes and quantum error cor-*

- rection. Ph.D. thesis, California Institute of Technology (1997). arXiv:quant-ph/9705052.
- [38] Kempe, J., Kobayashi, H., Matsumoto, K. & Vidick, T. Using entanglement in quantum multi-prover interactive proofs. *Computational Complexity* **18**, 273–307 (2009). Earlier version in Complexity’08, arXiv:0711.3715 [quant-ph].
  - [39] Cleve, R., Slofstra, W., Unger, F. & Upadhyay, S. Perfect parallel repetition theorem for quantum XOR proof systems. *Computational Complexity* **17**, 282–299 (2008). Earlier version in *Proc. 22nd CCC*, arXiv:quant-ph/0608146.
  - [40] Winter, A. Coding theorem and strong converse for quantum channels. *IEEE Trans. Inf. Theory* **45**, 2481–2485 (1999).
  - [41] Ogawa, T. & Nagaoka, H. Making good codes for classical-quantum channel coding via quantum hypothesis testing. *IEEE Trans. Inf. Theory* **53**, 2261–2266 (2007). arXiv:quant-ph/0208139.
  - [42] Magniez, F., Mayers, D., Mosca, M. & Ollivier, H. Self-testing of quantum circuits. In *Proc. 33rd ICALP*, 72–83 (2006). arXiv:quant-ph/0512111.
  - [43] Benedetti, R. & Risler, J.-J. *Real algebraic and semi-algebraic sets*. Actualités Mathématiques (Hermann, 1990).
  - [44] Nielsen, M. A. & Chuang, I. L. *Quantum computation and quantum information* (Cambridge University Press, Cambridge, 2000).
  - [45] Knill, E. Quantum computing with realistically noisy devices. *Nature* **434**, 39–44 (2005).

**Acknowledgements** We thank Edgar Bering, Anne Broadbent, André Chailloux, Matthias Christandl, Roger Colbeck, Tsuyoshi Ito, Robert König, Matthew McKague, Vidya Mahavean, Renato Renner, Shivaji Sondhi and Thomas Vidick for helpful conversations. Part of the work conducted while F.U. was at UC Berkeley, and B.R. at the Institute for Quantum Computing, University of Waterloo. B.R. acknowledges support from NSERC, ARO-DTO and Mitacs. U.V. acknowledges support from NSF grant CCF-0905626 and Templeton grant 21674.